



International Trade
and Forfeiting Association

itfa.org
info@itfa.org

FCC Initiative Best Practice Paper

Advanced Technology Operationalisation Framework for Financial Institutions & Vendors

October 6th 2025

by

Elizaveta Savinykh

Aurelia Bucicoiu

The views expressed by the authors in this publication do not represent their employers or any agency. Any content should not be interpreted in any capacity as regulatory guidance, recommendations, or interpretations of law. Such expressed views are to be considered individual opinions and perspectives.

Contents


EXECUTIVE SUMMARY	2
OVERVIEW	5
Context	5
Environmental Tension Points	5
Our Focus.....	8
FRAMEWORK	10
Guiding Principles.....	11
Joint Pursuit of Value	11
MRM Governance as a Strategic Enabler.....	12
Aligned Documentation	13
Hybrid Multidisciplinary Teams	14
Data & AI Literacy	15
Advanced Technology Deployment Lifecycle: Best Practice	17
Problem / Opportunity Identification & Solution Selection	18
Mobilisation	21
Implementation / Deployment	23
Production	25
CONCLUSION / CALL TO ACTION.....	27
Appendix	28
Appendix 1 – Documentation Standards.....	28
ML Model Card Example	28
Appendix 2 – Technology Impact Assessment Questionnaire.....	35
Appendix 3 – MRM / AI Vendor Requirements Overview	36

EXECUTIVE SUMMARY

Goal

This paper looks to outline a common framework for financial institutions (FIs) and vendors for the operationalisation of advanced technology, such as artificial intelligence (AI) and machine learning (ML), as well as sophisticated deterministic rules-based models. The proposed framework consists of five guiding principles with an overview of how these principles should be embedded in the advanced technology deployment lifecycle. The adoption of this framework by both FIs and vendors would lead to stronger model risk management (MRM) outcomes while also boosting vendor competitiveness, efficiency, and reputation and enabling FIs to strengthen governance, enhance transparency, and accelerate adoption of new technologies.

Framework Structure



The need for greater alignment between FIs and their vendors is driven by three key environmental tension points:

- **Nature of the Technology:** Advanced technologies are increasingly complex, often exhibiting reduced explainability and interpretability; meanwhile, existing FI control frameworks and processes may not be fully fit for purpose to effectively govern the technology without stifling its adoption.
- **Regulatory Change and Uncertainty:** Regulatory scrutiny around AI has been growing in recent years, largely prompted by the advent of Generative AI (GenAI). However, this is extending to broader technology areas, and regulators have expressed concerns.
- **Increased Reliance on Vendors:** As advanced technologies become more complex and specialised, reliance on the use of third-party solutions is growing while the knowledge gap between vendors and FIs is widening. This is exacerbated by IP considerations and the opaque nature of the vendor ecosystem.



To address these factors, the framework in this paper sets out five guiding principles and recommendations:

1. Joint Pursuit of Value: FIs and vendors should work in close, transparent alignment to define and track detailed success outcomes, spanning both core objectives and additional value areas. Regular value check-ins are recommended to reinforce long-term Return on Investment (ROI) and shared accountability.
2. MRM Governance as a Strategic Enabler: While FIs continue to maintain ultimate accountability for all deployed solutions (including all parameters, configurations, and controls) there should be a greater expectation for vendors not only to provide technology documentation but also to clearly highlight relevant risks, limitations, and constraints of their technologies and recommend appropriate mitigation and control measures.
3. Aligned Documentation: Vendors should adopt consistent documentation standards that support FI MRM needs. Recognising varied maturity levels, a sample "Vendor Documentation Card" is proposed as a baseline to improve transparency and ease of use across the FI industry.
4. Hybrid Multidisciplinary Teams: Contractual agreements should clearly delineate the support services that are included with product licenses (e.g., product documentation) versus those requiring a separate statement of work specific to the deployment (e.g., data quality reports). Depending on the nature of the engagement, vendors should commit to a certain level of availability from their internal technical and governance-focused experts.
5. Data & AI Literacy: Both FIs and vendors should foster organisation-wide data and AI literacy. These competencies should not be limited to technical teams but instead be integrated across the broader workforce to build a culture of informed oversight and collaboration.

The final section of the framework focuses on the four lifecycle phases of advanced technology deployment, embedding the above principles, and outlining a structured impact assessment – including a practical questionnaire – to help evaluate technology outcomes, prevent unintended consequences, and ensure alignment with defined value objectives:

1. Problem / Opportunity Identification & Solution Selection
2. Mobilisation
3. Implementation / Deployment
4. Production



Call to Action

As FIs continue to adopt increasingly complex technologies and deepen their reliance on third-party providers, stronger FI/vendor alignment becomes critical.

To support adoption, we recommend that:

- The FI industry work toward standardising MRM requirements to reduce friction and enable more consistent vendor engagement. This process could be greatly supported by the buy-in and endorsement of relevant industry bodies.
- Vendors align with FI expectations by developing standardised documentation and support models based on the proposed framework.
- Over time, vendors should consider embedding MRM capabilities directly into their solutions, enabling proactive governance and innovation, such as MRM-native tooling and agents.



OVERVIEW

Context

In recent years, the pace of technological innovation has accelerated significantly – particularly in the areas of artificial intelligence (AI) and advanced analytics with the arrival of generative AI (GenAI). As a result, financial institutions (FIs) now benefit from sizable opportunities but are also exposed to considerable risks. Many FIs are under growing internal pressure from boards and executive leadership to rapidly unlock value from these new technologies. Yet while the urgency to avoid falling behind competition is real, it is counterbalanced by significant regulatory uncertainty.

Traditional control frameworks, skillsets, and governance structures are likely to prove inadequate or too operationally burdensome when applied to innovative initiatives, thus both increasing risk and slowing down adoption. The share of AI and complex technology pilots is rising; however, many remain at the proof of concept (PoC) stage without progressing to full operationalisation. According to IBM's CEO study¹, surveyed CEOs report that only 16% of AI initiatives have scaled enterprise wide. And when these technologies are deployed into production but with improper understanding and management, they can introduce new vulnerabilities into control environments and fall short of delivering their expected value. This risk is not escaping regulatory attention.

Environmental Tension Points

There are three key tension points that FIs must consider:

1. Nature of the Technology

The use of advanced technologies (which include AI and machine learning (ML) as a core subset of AI, as well as sophisticated deterministic rules-based models) is not new in the banking sector.

However, due to the latest advancements in technology and data processing, the nature of available solutions is shifting towards greater inherent complexity, and this is frequently obscure. Meanwhile, the significant focus and investment kickstarted by the advent of GenAI in recent years is not only prompting FIs to invest in GenAI², but also to consider other advanced technologies in more use cases. Some of these technologies have already been available for several years, but there is now a tangible drive to pursue their implementation in areas where simpler processes may have historically prevailed.

This focus on AI and other advanced technologies is putting additional pressure on FI governance frameworks, particularly in relation to model risk. When model risk was originally defined and embedded within FI risk management, it was largely concerned with statistical models built on assumptions and probabilities. Although complex, these models allowed for a level of

¹ [IBM Study: CEOs Double Down on AI While Navigating Enterprise Hurdles](#)

² ["Per the 2024 Gartner CIO and Technology Survey, 42% of banking CIOs have deployed, or are planning to deploy, generative AI \(GenAI\) in the next 12 months"](#)



explainability and interpretability: their steps could be understood, and their outcomes reproduced. They were principally designed for situations with uncensored data, where adverse outcomes would usually be visible.

In contrast, AI models represent a conceptual shift. Their input is data, but the core process that produces the model, machine learning, often does so in ways that are opaque or at least difficult for a human to grasp. The latest (typically nondeterministic) models, such as GenAI, present an even greater challenge, both in terms of accuracy and in terms of explainability and interpretability: there is limited guarantee that the same results can be reproduced given the same set of inputs and parameters.

As a result, the controls and processes that are effective at governing traditional and less complex models may prove to be unsuitable for more advanced technologies, both in terms of design and scope. It is important to stress that this is not only true for GenAI, but also for the broader universe of advanced technologies.

2. Regulatory Change and Uncertainty

Much like the 2008 financial crisis highlighted the need for robust model oversight – accelerating the definition and development of Model Risk Management (MRM) frameworks – the rise of GenAI is now prompting a major shift in regulatory thinking and guidance.

MRM originally emerged as a discipline to address the need for effective oversight of regulatory models. The U.S. Federal Reserve was the first to issue comprehensive supervisory guidance on MRM in 2011³, with other jurisdictions – such as the EU, UK, and UAE – subsequently introducing their own standalone consolidated guidelines.

Over time, the scope of MRM has expanded to include internal models used for areas like financial crime risk management and operational risk. While AI is not a new concept, regulatory bodies around the globe are acknowledging that GenAI introduces new challenges that existing frameworks may not be fully equipped to handle. As an example, in a recent survey of AI usage in the local financial services industry, Luxembourg authorities highlighted that only 56% of use cases reported good (25%) or very good (31%) auditability⁴, representing a significant decrease vs. the previous survey. While the reasons for the downgrade cannot be easily explained, it is speculated that it may well be due to “the increasing level of complexity of the AI solutions used and the difficulty in auditing them, together with more realistic scores provided by respondents based on more experience”.

Concerns around existing frameworks being fit for purpose for GenAI are prompting jurisdictions that do not have standalone MRM guidelines to nonetheless issue dedicated AI guidance and requirements. The EU AI Act⁵ is notable for being the world’s first comprehensive AI law, but other jurisdictions are actively developing their respective approaches.

³ [The Fed - Supervisory Letter SR 11-7 on guidance on Model Risk Management -- April 4, 2011](#)

⁴ [Second thematic review on the use of Artificial Intelligence in the Luxembourg financial sector – CSSF](#)

⁵ [EU AI Act: first regulation on artificial intelligence | Topics | European Parliament](#)



However, despite a growing number of publications and industry consultations on the topic, regulatory uncertainty persists. Definitions of AI differ (although most major publications are converging on the definition adopted by the Organisation for Economic Cooperation and Development (OECD)⁶) and are subject to interpretation from technically narrow to very broad, impacting which models might fall into scope. Differing global regulatory requirements and sophistication introduce the risk of regulatory arbitrage, although FIs with large global footprints are likely to adopt the strictest standards and thus face the highest costs. Furthermore, while much of the public focus is on AI, all types of advanced technologies are likely to come under scrutiny due to the overall sharpened regulatory attention. As an example, in its 2025 Opinion⁷ the European Banking Authority has highlighted that “RegTech solutions offer significant potential for better compliance and a reduction of manual errors, but their successful deployment has been hampered by inadequate in-house expertise, poor governance and insufficient oversight.” This reference is not limited to AI solutions.

Lastly, the role of vendors in MRM remains ambiguous. Existing guidelines acknowledge vendor-related challenges and set out vendor-related requirements for FIs from which expectations for vendors can be inferred (refer to *Appendix 3 - MRM / AI Vendor Requirements Overview* for detail). However, the requirements are not consistent and do not always outline the detail of the expected vendor contribution. Greater and more explicit clarity in this area could enhance MRM’s applicability and standardisation across the broader FI/vendor ecosystem.

3. Increased Reliance on Vendors

Paradoxically, while access and awareness of advanced technologies is becoming more democratised – this is true across several technologies thanks to the focus generated by GenAI – their complexity means that FIs are increasingly unlikely to successfully develop such models internally. The resources, expertise, and data required are simply too great. This would not be cost-effective even for the largest FIs to do, particularly in use cases that sit outside of the FI’s core business. Further, FIs frequently turn to vendors to solve complex problems that cannot be addressed quickly or effectively in-house. Consequently, reliance on third-party vendors is set to grow substantially across all solution types – both standalone and integrated into proprietary internal solutions.

In addition, there is increasing inter-reliance amongst the vendors themselves. In the GenAI space, only a small number of the largest vendors have access to the computing power and training datasets needed to develop large-scale models such as large language models (LLMs) (a priority focus for FIs in the GenAI space). As vendors build solutions with easier integration in mind (e.g. via APIs), FIs have the opportunity to integrate multiple vendor solutions at once, bringing in progressively narrow expert applications. As a result, vendors and FIs increasingly

⁶ [AI Principles Overview - OECD.AI](#): An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.

⁷ [A careless use of innovative compliance products can lead to money laundering and terrorism financing risks, the EBA says in its Opinion | European Banking Authority](#)



integrate or build on other vendors' solutions, compounding complexity and creating further layers of reliance.

This reliance introduces new challenges. The knowledge gap between FIs and their vendors is widened not only by the complexity of the technologies but also by intellectual property (IP) protections and an understandable lack of internal expertise on each vendor solution. The vendor ecosystem itself is highly varied and can be opaque, further complicating risk assessments, internal governance, and external oversight. Finally, while industry and regulatory definitions of AI systems and the understanding of standards are converging, the definitions and standards used in the vendor and vendor analyst space continue to vary significantly, further reducing transparency.

Our Focus

Ultimately, FIs remain responsible and accountable for managing model risk within their business. However, as reliance on third-party vendors grows, there is a clear opportunity to reduce the resulting MRM challenges by developing a set of best-practice recommendations to guide FI-vendor interactions and expectations in the context of MRM.

Vendors possess deep technical expertise in their own solutions but generally have limited MRM focus. On the other hand, while FIs bring strong MRM and risk management capabilities, they may lack in-depth knowledge of the relevant aspects of the third-party technologies they employ.

The objective of this paper is to propose a common framework for FI-vendor operationalisation standards consisting of guiding principles and best practices across the advanced technology lifecycle.

Definitions:

- *For the purposes of this paper, the definition of advanced technology includes AI and ML (as a subset of AI) but also sophisticated deterministic rules-based models, as, despite their non-probabilistic nature, their complexity can present similar governance challenges.*

Scope:

- *This paper applies to all regulated FIs with a strong focus on banks and all vendors providing solutions to such FIs.*
- *This paper does not consider how advanced technology could be used to optimise or enhance the risk management process of other advanced technologies (e.g., MRM agents); however, all principles outlined in this paper would be equally applicable to such a process.*
- *This paper does not consider systemic industry risks associated with outsourcing (e.g., industry-wide single-vendor dependencies).*

Benefits for Vendors and FIs

While vendors do not operate under the same oversight regimes as FIs, they nonetheless stand to benefit considerably from the adoption of a shared and industry accepted MRM FI-vendor structure:


- Increased win probability: The adoption of clear and industry-recognised MRM standards increases a vendor's competitiveness by aligning with FI requirements and expectations from the outset. Failure to demonstrate compliance with MRM practices can, at worst, disqualify a vendor from consideration entirely or, at best, reduce credibility during procurement processes and materially impact the chances of winning new business.
- Reduced delivery cost: Standardised expectations streamline documentation processes, helping vendors reduce rework and duplication, particularly when engaging with multiple customers across jurisdictions.
- Internal capability building: Adopting a consistent MRM framework also supports internal capability development, enabling teams to build lasting expertise in model lifecycle-management best practices, such as robust documentation and testing. This not only results in a more resilient product offering, but also lays a solid foundation for responsible innovation. These benefits can be amplified by embedding MRM capabilities directly into vendor solutions, delivering MRM by design with inbuilt transparency and governance from the outset.
- Reputational impact: Vendors who embrace strong MRM practices can position themselves as leaders in responsible innovation.

FIs, currently fully responsible for MRM, would similarly benefit from a common framework:

- Stronger MRM controls: Consistent vendor standards support more robust governance by making it easier to assess compliance, conduct model monitoring, identify control gaps, and uniformly enforce controls across all vendors.
- Improved vendor transparency: The use of a common framework enables clearer early alignment between vendor solutions and FI risk appetites. This improves visibility during procurement processes, simplifies vendor comparisons, and helps inform the selection of acceptable technology solutions.
- Faster operationalisation of new technologies: A shared framework reduces redundant evaluation steps, allowing FIs to more quickly adopt and integrate vendor solutions.
- Streamlined MRM processes across the organisation: Standardised templates and expectations reduce the effort to request, customise, and rewrite documentation in support of MRM activities. This results in time savings across MRM, procurement, technical and business functions.



FRAMEWORK





Guiding Principles

Joint Pursuit of Value

Value realisation should be the focal point of any complex technology initiative. While each FI has unique pain points and value drivers, vendors bring deep expertise in their technologies, informed by market research, prior implementations, and a clear understanding of constraints and opportunities of their solutions.

Technology success outcomes are the long-term value deliverables that can be achieved once a technology is fully live and operational. Vendors can help shape these outcomes by advising on how their solutions can most appropriately be configured and scaled to address the FI's goals. Further, they can validate the scope and ambition of the outcomes based on prior market experience and the constraints of the project to ensure the outcomes are achievable. Vendors can support the definition, embedding, and long-term outcome ownership by the FIs through stakeholder education on their technology, including broader capabilities and limitations. Lastly, depending on their experience, vendors may be able to provide early guidance on broader operationalisation requirements beyond the technology alone (such as processes, education, feedback loops, skillsets, etc.) to make their solutions successful. In addition to the primary FI objectives for the project, vendors can help identify other areas of impact by contributing value trees to map out further cost-effective benefit streams specific to their technology.

Nonetheless, it is important to note that the achievement of success outcomes is a joint endeavour, and does not fall solely to the vendor. A number of FI decisions and workstreams, from scope definition to delivery phasing to the operationalisation of the technology, can have a substantial impact on overall success, and such constraints should be duly acknowledged by the FIs.

Following success outcome identification, vendors are also well placed to outline the best approaches for how to measure against them – specifically by setting out the available data points, and by suggesting alternative measurement approaches where needed.

Finally, vendors can support long-term tracking and communication of progress against success outcomes via periodic value check-ins. These can include a broad range of activities, such as joint measurements of the total cost of ownership (TCO) vs. achieved benefit, robust upgrade planning, overviews of new features and roadmap items, and the surfacing of new requirements and co-innovation opportunities. Similarly, should the technology fail to deliver to its expected outcomes, vendors can assist with root cause analysis and take steps to close the gaps. These types of collaborative activities help ensure a sustained return on investment (ROI) for the technology and reduce the risk that the solutions will degrade over time due to lack of focus and investment.



Our recommendation is a close, transparent alignment between FIs and vendors in defining and tracking detailed technology success outcomes across both primary objectives and additional value areas. This process should be supported by periodic FI/vendor value check-in sessions aimed at driving long-term ROI commitment.

MRM Governance as a Strategic Enabler

MRM should not be viewed merely as a compliance obligation, but as a foundational enabler for successful technology adoption. When integrated holistically across the project and model lifecycle, MRM can facilitate value realisation, ensure progress beyond PoCs, and reduce barriers to full adoption. If a model is designed from the outset with risk management at its core, there will be less friction as it progresses through the relevant assurance processes. While FIs remain fully responsible and accountable for MRM within their business, vendors can play a more defined role in helping to shape robust controls around their technology.

FIs should aim to define a clear risk appetite for advanced technologies based on the acceptable levels of explainability and interpretability for each use case type. This will streamline early vendor selection by clearly defining what types of technologies and risks are acceptable for the target purpose. A risk-tiered framework, which applies differentiated assurance standards based on model materiality and complexity, can enable more agile experimentation in low-risk environments while safeguarding critical areas of the business. However, as part of this process, FIs should remain cautious of inconsistencies in how vendors define AI, and should review these in the context of internal governance definitions. Vendors, in turn, should support FIs in understanding and assessing the inherent complexity of their technologies.

Traditionally, MRM as a process is fully owned and driven by FIs, with vendors playing a supporting role. FIs are responsible for articulating MRM frameworks, specifying and designing controls, while vendors primarily contribute by providing solution documentation. This is closely aligned to regulatory and industry expectations.

However, as technology continues to evolve in complexity, it may no longer be feasible for all FIs to sustain extensive internal expertise across every emerging technology domain. The appetite for building such expertise will depend on the FI's size, technological maturity, and business model, as well as the level of specialist knowledge required by the technology. Regardless, the cost of such expertise will continue to rise. This shift therefore necessitates a more collaborative model between FIs and vendors.

While FIs will continue to own the risk and maintain accountability for the implemented controls, there should now be a greater expectation for vendors to proactively highlight applicable risks specific to their technology and to suggest appropriate control options for mitigation, monitoring,



and measurement. These options should cater for varying FI risk appetites (e.g., from more expensive but holistic control options to lighter-touch health-check sampling controls). Vendors, given their deeper familiarity with the intricacies of their own solutions, are often better positioned to define a reasonable starting point for how their technology should be risk-assured.

This type of vendor MRM engagement is only possible if vendors achieve a sufficient level of internal expertise in MRM to oversee their own model development processes, especially when they are providing fully built models or solutions as a service (SaaS). The more complex the technology – including in terms of the integration of other vendor solutions into the final product – the deeper the required level of MRM expertise to manage key risk areas such as fairness, explainability, interpretability, accountability, transparency, and reliability.

That said, FIs remain responsible for understanding, evaluating, and challenging the articulated risks and proposed controls to ensure they are fit for purpose. Ultimate decisions on all parameters, thresholds, and configurations remain with the FIs, even where these decisions are informed by vendor input. Vendors, in turn, are responsible for ensuring FIs have a sufficient level of information to understand the risks and to challenge the controls and configurations. This evolving partnership cannot extend to operational risks and controls beyond the technology itself – such controls fall outside the vendor's scope, and remain solely within the FI's domain.

It is worth noting that while many organisations are adopting AI strategies, there is not necessarily a need for them to develop entirely new AI governance structures. Instead, existing MRM frameworks can be reviewed and expanded to address the heightened complexity and the additional risks associated with AI. In simple terms, effective MRM governance is AI governance.

Our recommendation is a closer alignment between FIs and vendors on the design of appropriate control frameworks for advanced technology solutions. In addition to providing technology documentation, vendors should clearly highlight the relevant risks, limitations, and constraints of their technologies, and recommend appropriate mitigation and control measures.

Aligned Documentation

Aligning technology industry documentation standards with those of regulated industries reduces the friction associated with the knowledge transfer from vendors to users and increases the transparency and explainability required for a successful deployment.

MRM requirements mandate that FIs have robust documentation standards for their models and that they include a clear description of what should form part of good model documentation. Conversely, vendors don't have any specific guidance or regulation to indicate the level of documentation they must produce. However, MRM also requires that FIs ensure a sufficient level



of detail in third-party vendor model documentation to validate their use of the model. Therefore, part of the regulatory requirements can be extrapolated to vendors.

From a vendor perspective, good documentation can support their commitment to transparency and decrease the total cost of ownership for their users, which can be an important competitive advantage. It is therefore recommended to align to the industry framework, called *model cards*, to encourage transparent model reporting. Model cards are an idea originally explored in a Google research paper in 2018⁸, and have been used across the industry to organise essential facts of machine learning models in a structured way. Since model cards were first proposed by Mitchell et al. in 2018, they have been adopted and adapted by various organisations, including by major technology companies and startups developing and hosting machine learning models, researchers describing new techniques, and government stakeholders evaluating models for different projects⁹.

However, industry-standard model cards only target machine learning models. FIs may wish to apply the MRM framework to decision-based rules or algorithms that are not classified as models, but which are complex in nature and have a material bearing on business decisions¹⁰. Also, while model cards include most of the required information for MRM model documentation, they are not structured in a way that supports a smooth mapping of information between the vendor and FI MRM standards. Creating a hybrid between the standards of the technology industry and FI MRM requirements would decrease the friction between the vendor and the FI when leveraging information from the vendor. Common standards would make producing internal model documentation simpler.

Our recommendation is for vendors to adopt a documentation standard that considers the MRM needs of their users. Acknowledging vendors have different maturity levels, a proposed template for a Vendor Documentation Card is provided in Appendix 1 – Documentation Standards as a starting point.

Hybrid Multidisciplinary Teams

Setting up teams that represent all areas of the organisation impacted by the advanced technology – while embedding vendor expertise – ensures knowledge synergies and protects against blind spots.

The success of an advanced technology deployment relies on the involvement of the right people with the right skills and knowledge; this can be achieved via multidisciplinary teams. The involvement of the selected vendor (or a professional services partner with a similar level of

⁸ <https://arxiv.org/abs/1810.03993>

⁹ <https://huggingface.co/docs/hub/en/model-card-landscape-analysis>

¹⁰ <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2023/ss123.pdf>



knowledge) is critical given the required depth of knowledge of the deployed technology. Depending on the level of collaboration between the FI and the vendor, this step could be at least partially included as part of the selection process to avoid situations in which IP protection or unclear contractual terms and conditions limit access to the required information. This is especially important in cases where the discussions might result in the decision to discontinue the deployment.

Once an FI made the decision of deploying advanced technology developed by a vendor, it is important to identify the different roles from both organisations that must be involved in the process, and their different levels of participation during the different stages of the lifecycle. Depending on the size and maturity of the FI, some of these roles might be covered by companies providing professional services, and in some cases, even by the vendor that provides the technology – just from a different perspective.

Advanced technology deployments require traditional roles like data engineers, data scientists, architects, and business analysts, but there are also some nascent roles that organisations should consider for successful implementations. From an MRM perspective, the role of a model manager is particularly important (especially for ML models) to ensure the model is set up correctly and that the processes around it behave as expected throughout its life cycle, including drift monitoring and the selection of existing (foundation) models for (re)use. While embedding technology ethics across the organisation, an AI ethicist could ensure a structured approach to considering the unintended consequences of the use of data and AI, and to determine how to best manage risks and opportunities.

Our recommendation is to clearly state in the contractual agreement which types of support are included with the product license (e.g., product documentation) and which are subject to a separate statement of work specific to the deployment (e.g., data quality reports). It is important to ensure the vendor commits to a certain level of availability from their internal resources responsible for product development (engineers or data scientists) and for AI governance.

Data & AI Literacy

Treating data literacy and AI literacy as core competencies across the workforce and allocating significant leadership attention to them ensures knowledge is disseminated across every role that requires an understanding of how technology impacts their specific function.

Having a multidisciplinary team with the right roles is not sufficient if those roles lack the data/AI literacy required to deploy advanced technology. AI literacy is the ability to effectively and responsibly utilise AI in a business and societal context with competency to identify relevant use



cases, as well as to implement and operate corresponding AI applications. Data literacy entails enabling employees to consume, analyse, and make informed decisions with data.

Most companies have focused mainly on training, attracting, and retaining the next generation of AI scientists, and less on ensuring data/AI literacy across their wider workforce and SMEs, which would enable them to competently engage with the AI/technology workforce. This lack of AI skills and organisation-wide AI literacy stands in the way of scalable and responsible technology operationalisation. If managers lack AI literacy, they may fail to foresee the shifts in organisational structure, business processes, and culture needed to adopt AI solutions, or be unprepared to implement those changes.

The challenges companies face with operationalising advanced technology are not new, but the pace of AI development, especially GenAI, has created added pressure. FIs have been deploying technology for decades to support their objectives, be they statistical models or machine learning models focused on the analytical side of AI. This has required them to build their own AI talent pool, mainly in the areas of quantitative analysts, data analysts and data scientists, software developers, and so on. Because most of the advanced technology fits the definition of a model as described in MRM regulatory guidance, FIs have also worked to ensure their compliance specialists are familiar with the technology behind the models. This was easier to do when those models were quantitative methods that applied statistical theories already familiar to most parties involved in their deployment. However, advances in technology and data processing power have permitted not just more complex deterministic quantitative methods (such as decision-based rules or algorithms), but also stochastic and dynamic (and therefore nondeterministic) systems like GenAI. This complexity makes it extremely difficult for the typical roles involved in technology deployments to understand to the same level the technology they are deploying.

Our recommendation is for both FIs and vendors to consider data literacy and AI literacy as core competencies across the workforce, rather than implementing data/AI literacy programmes solely for those in highly technical roles.



Advanced Technology Deployment Lifecycle: Best Practice

Deploying advanced technology offers opportunities, but it also entails risk. Because of this, throughout the technology lifecycle it is important to have clarity on the impact of the deployed technology to prevent unintended negative consequences and ensure value objectives are met. This can be done via an impact assessment, preferably based on a clear list of questions that support the thought process and increase the accountability, quality, and reproducibility of the deployment. The best practices listed below are synthesised in a questionnaire that could support such an assessment. Refer to *Appendix 2 – Technology Impact Assessment Questionnaire* for the consolidated list of questions.

The below framework provides best-practice considerations throughout the different technology lifecycle phases.

Problem / Opportunity Identification & Solution Selection	Initially the framework should be used in this stage to investigate if the use of advanced technology is feasible and desirable. Key questions could be included in the vendor RFI / RFP process to ensure early involvement of potential vendors. Once a vendor is selected and determined to be in line with the initial findings and actions, the framework should equally inform legal / contractual discussions to guarantee an appropriate level of vendor support during the deployment.
Mobilisation	The observations from applying the framework in the pre-deployment phase should be used in this stage as a to-do list, with follow-up activities included in the project plan.
Implementation / Deployment	During this stage, the framework should be used as a checklist to ensure relevant aspects have been taken into consideration, either by the vendor or the FI, as per contractual agreements.
Production	Once the technology makes it into this stage, it's important to ensure that any of the changes resulting from upgrades to the technology or its uses continue to conform to the initial requirements.



Problem / Opportunity Identification & Solution Selection

Due to the risks inherent in advanced technologies, the question of whether a technology's impact is proportionate to the intended objective is critical to consider at the very beginning. While AI presents many promising opportunities, it is important to avoid adopting the technology solely to follow trends or capitalise on the current hype; the application of AI should serve a clear, strategic purpose rather than becoming an end in itself. FIs should clearly identify the intended purpose and results of the technology and review potential solutions with multiple vendors.

Each vendor should be able to provide sufficient information (i.e., enough to support the main solution selection with regards to compliance and technology management teams) for the FI to decide which solution suits their needs better, and also to provide reasoning for the final choice:

- Intended purpose and intended result of model → Vendor documentation should include a description of the use case to which the technology is applied, the business process in which it is embedded, and the potential key objectives (which can be fully or only partially aligned to the objectives of the FI). FIs should request that vendors include success outcomes on measuring these key objectives based on their experience with their users, including pragmatic methodologies that ensure it is feasible to quantify these metrics; this is critical to forming an objective assessment of the deployment's success. They should be explicit about any out-of-scope uses or scope limitations for the technology (e.g., high performance achieved only on one language, or it cannot be used for certain jurisdictions due to data availability), and vendors can recommend controls that could be used for the listed limitations (e.g., manual controls on a sample dataset).
- Description of the intended solution → Vendor documentation should provide a description of the technology used (e.g., a rule-based system, a machine-learning model, a large language model (LLM)) and, if considered, indications on other alternative theories or approaches, including the reasoning behind the selected technology (e.g., the performance of the ML model was substantially higher compared to the rule-based model, justifying the additional risk of lower transparency and explainability). Once the FI selects one vendor over another, they should include if the selection was done (also) because of technology differences. Is important to be explicit about all the technologies used in the deployment, including some tools that are only auxiliary but could be considered as feeder models by some MRM frameworks (e.g., using an LLM to identify data quality issues might be unacceptable for an FI, due to the risk related to including data in LLM prompts).
- Prerequisites for the intended solution → High-quality and sufficiently voluminous data are essential to create, test, evaluate, and validate models. The success of the deployment relies on data with sufficient quantity, quality, cleanness, and structure, not just to train a model, but also to use it for inference purposes. Vendors need to be clear on the required characteristics for the input data to achieve the best output from the model. Resolving data challenges is a priority for any advanced technology project, because incomplete, inaccurate, or disconnected data will negatively impact the output of even the best trained



model. By some estimates, nearly 80% of AI projects fail.¹¹ Without high-quality, trusted data, training and deploying complex technology models is unlikely to meet success outcomes.

- Maturity of the features included in the intended solution → As expected with fast developing technologies, the deployed solution may also include emerging features that bring more value more quickly but that may also pose more potential risk. The FI needs to balance the performance gains with the additional risks, such as: lower levels of market proof of efficacy of an approach; less support or weaker documentation provided by default by the vendor; and even the possibility that the feature might be depreciated as a result of development challenges or a shift in priorities. These conditions should be explicit so that the assessment of the impact of the technology is accurate and proportionate to the intended objectives and performance gains, as emerging features might be more suitable for some purposes than for others.
- Alternatives to the intended solution → As part of technology development, vendors typically explore alternative approaches that could achieve the same objective and therefore can provide an overview of potential benefits of their technology vs. these alternatives. The FI should enhance this information with additional data collected from market analysis and indicate if no technology, less complex technology, or a different type of technology could be used for the same objective and articulate reasoning for the selection. This reasoning should consider not only the benefits but also the risks in terms of transparency, explainability, bias and potential costs associated with them.
- Type of deployment for the intended solution → Vendors should discuss with FIs how their product is deployed to ensure it is properly integrated in the existing architecture and to avoid difficulties in maintenance post-deployment. Usually this also provides an idea on the level of control the FI will have of the solution and the split between the changes that could be done only by the vendor or by the user. For example, an ML model can be deployed in a way that does not permit the user to retrain the model, so any notice of model drift could only be handled by the vendor.
- Licensing attached to the intended solution → This is not a new dimension, but one that has been made more visible by the developments in LLMs, which are often built on vast amounts of data sourced from diverse origins. This raises questions about ownership, data access, and intellectual property. The legal landscape surrounding AI is inherently complex and legislation struggles to keep pace with these advancements, as traditional frameworks for intellectual property and data protection were not designed with AI in mind. Moreover, different jurisdictions may impose different standards, and this patchwork of regulations complicates the global deployment of AI technologies. Also, there is a growing acceptance that the power of an AI model is strongly related to the quality and volume of data used to train it. Given the fact that many vendors own the algorithms but not the data, the most powerful solutions are often those with the most creative approaches to obtaining data. It is therefore critical that FIs are able to obtain from

¹¹ [The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI | RAND](#)



vendors an appropriate view of the data used to build solutions, that FIs have an up-to-date understanding of the licensing implications of different data choices, and that there are clear agreements on the IP rights for all data involved in the deployment, including non-production data (such as testing data).

- Compliance with laws and regulations → Vendors should be able to indicate the industry standards, laws, and regulations to which they comply, but ultimately it is the FI's responsibility to ensure (and be able to prove) that all necessary requirements have been implemented by the vendor or themselves. While at this stage it is not critical to collect evidence for everything, it is important to ensure each party is aware of its responsibilities and commits to further engagement where needed. For example, the vendor can confirm that it has retained the data used for training the model and that it has a process in place to share it with users, if needed; the FI does not need to ask for the data itself, unless required to do so by a regulator.

Most of the above information is also included in the general MRM practices, and will be required for model validation purposes as deployment progresses. Collecting this information from the vendors upfront serves also as a check on their ability to support the MRM activities down the line. At this stage, FIs should only aim to request the information at a higher level and from short-listed vendors to avoid additional unnecessary effort.

Successfully deploying advanced technology requires a strong combination of roles and skills, on the part of both the vendor and the FI. The intended solution is not just about the technology, but also about the people who configure it, deploy it, maintain it, and use it. This is where the vendors should be able to clearly indicate which roles and support are included as part of the product, and which roles can be performed by either the vendor, the FI, or a third party (e.g., professional services partners). Mature FIs that have available resources with the correct skill sets might want to limit the additional cost that comes from professional services (be they from the vendors or third parties). In this case they need to be even more certain that their contractual agreement covers the knowledge transfer that is specific to the technology; this cannot be covered only by the technical documentation. In general, more mature FIs also have a higher need for customisation, and therefore it is critical to have an exact understanding of the impact any given parameter can have on the model, as well as methods to measure and track this impact. In some cases, technology products have a faster pace of development, which results in a quicker release cadence than for traditional IT. For FIs that have relied on vendor or third-party resources to deploy the system, a lack of accumulated information on how to do both a technical and functional upgrade can become a challenge, and can result in higher costs or missed value.

The considerations listed above should be used, in conjunction with measurable functional and technical requirements and market referenceability¹², both to select the best solution and to prepare the stage for the deployment. To ensure that all agreements and clarifications made in the RFP phase are made explicit, both parties are encouraged to incorporate them into contractual agreements. Vendors should make sure their general Terms & Conditions include an AI

¹² Defined as being able to provide evidence of existing use or deployment in relation to vendor solutions



addendum that covers the split of responsibilities in regard to all of the following: compliance to laws and regulations, technical and functional documentation, appropriate data quality and relevance, appropriate testing, additional controls required by the system's limitations, performance monitoring, and model validation.

PROBLEM / OPPORTUNITY IDENTIFICATION & SOLUTION SELECTION - QUESTIONNAIRE

1. What is the intended purpose and result of the technology?
2. What is the intended solution to address stated purpose and result?
3. What are the prerequisites for the intended solution?
4. What is the overall maturity of the features included in the intended solution?
5. Are there any alternative solutions to address the described purpose and result?
6. How will the intended solution be deployed?
7. What are the licenses required for the use of the solution, and what are the licenses linked to the assets used in the development?
8. How is compliance with applicable laws and regulation being achieved?

Mobilisation

Following vendor selection, the analysis performed during the previous stage should be used to inform the advanced technology project plan for the deployment and ensure all relevant aspects are considered at the right time by the right people. Allocating sufficient time to deal with the complexity of deploying advanced technology increases the success of its operationalisation. Planning should not only be focused on the activities to be performed, but also on the resources that have the skills required to perform them. The conclusions might be that some roles need to be covered by the vendor, or, if this is not part of the agreement, by an independent third party, if that party has the required product knowledge.

Considering the following factors while planning projects with the vendor or other professional services companies helps set the right stage for the deployment:

- Risk tier allocated to the deployed model → The risk assessment of a model may consider factors such as materiality, complexity, and usage of the model. Since it is linked to its actual use, it is the FI's responsibility, but the vendor should offer information to be used as input for this assessment. Vendors should support FIs, especially when assessing a model's complexity, considering the risk factors that impact a model's inherent risk within each component of the modelling process. These include the nature and quality of the input data, the choice of methodology (including assumptions), the requirements and integrity of implementation, and the frequency and extensiveness of model use. It is obvious that more attention should be given to advanced technology that presents a high risk due to its complexity. However, doing so might result in unintended gaps caused by misinterpretation of the link between risk and complexity. For example, a low-complexity rule-based system can present high risks due to the lack of transparency from the vendor



side or due to the low level of interpretability on the user side. Considering this, our recommendation is to apply the assessment for any type of technology, but to adjust its depth based on its complexity and possible level of transparency and interpretability. To ensure more comparable inputs, it is recommended that vendors assess complexity using industry benchmarks provided by research and advisory firms.

- Potential risks and limitations associated with the deployed technology → Because vendors are the ones with the complete product knowledge, they are responsible for indicating the potential risks associated with their products, accounting for the recommended uses. While developing and deploying the product with different users, vendors can identify many foreseeable harms, misunderstandings, and technical and sociotechnical limitations, and therefore are able to provide information on warnings and potential mitigations. These mitigations can be translated by vendors into potential controls that the FIs must consider. The risks provided by the vendor are only technology-related, not business operations related. The risk controls must take into account the FI's risk appetite, but primarily stem from the product's limitations, which have been identified and accepted by the vendor throughout the product development cycle. Following identification of such product limitations, vendors are expected to provide options to address or remediate the concerns (e.g., through monitoring, by targeting only eligible population in the model deployment, by ensuring accuracy of labelled datasets for future validation workflows, etc.), and this information is key for the proper design of FI controls or for an informed risk acceptance by the FI.
- Specialised skills required for the roles involved in the activities → Ideally these skills have already been identified during the initial phase, and the current stage is about matching them to available resources within the FI, the vendor, or another party. Generally, stakeholders need to consider AI governance as a pillar of the deployment and safeguard the required resources and activities.
- Operating model required to drive value → Managers often struggle to understand how advanced technology/AI can address real problems in the workplace and simultaneously underestimate the enterprise-wide implications and changes in business culture that it may entail. Developing the right operating model to embed the technology in the business operations is very complex, and highly personal to each FI. Still, vendors can directly support with input for procedures, recommendations for process changes that support new or different controls, and structural feedback loops that ensure the rapid flow of information in both directions. Ensuring all these topics are captured in the project plan is a confirmation that the FI does not treat the implementation only as a technology deployment, but also as the required business transformation that it is.

This is also the moment in which FIs should initiate MRM framework activities to ensure the project team is aware of the model risk and governance-related tasks that need to be covered at different stages of the deployment. Involving a representative of the model validation team early in the process will ensure they can guide the work in such a way that their colleagues will have all the information required in the structure they expect to assess later in the final version of the model, when it is submitted for model validation. Including a model validation stream in the project

plan ensures that important governance activities are not missed and are not only run under pressure shortly before the go-live date. Similarly – and depending on the scale, materiality and novelty of the project – this may be a reasonable point to involve third line of defence representation, both to benefit from their input during the course of the project and to build relevant technology expertise in the audit function.

MOBILISATION - QUESTIONNAIRE

9. What is the risk tier allocated to the model, and what are the associated considerations?
10. What are the potential risks and limitations associated with the targeted technology?
11. What are the specialised skills required, and how are they sourced? If an external party is responsible for the deployment, what contractual agreements are in place?
12. What needs to change in the operating model to drive the targeted value outcomes?

Implementation / Deployment

While in the initial stages of the lifecycle it is recommended to collect information at a high level only, the implementation phase includes more detail, and in many cases requires greater knowledge about the technology and its impact against documented expectations and success outcomes. The target of this phase is to successfully deploy the technology into a production environment, which includes obtaining official approval from the model validation team.

Covering the areas below at the appropriate level of depth will set the groundwork for a successful run in production for the deployed technology:

- Transparency of the technology deployed → The vendor should educate the FI on the design, theory, and logic of the model, including the key steps from the modelling process and the algorithms used. Clear explanations should be included, where possible, of the degree to which underlying input-output relationships predict model outcomes. It is important to evaluate the applicability of selected algorithms and theories against the business objective and technology use, which should come as a more detailed view of the similar exercise run in the initial phase. Advanced technology presents different levels of explainability, and in many cases that is linked to the accuracy levels, so the vendor should indicate their reasoning behind any design decisions that trade accuracy for explainability. In some cases, a lower accuracy but higher explainability will fit better with the use case and the risk appetite of the FI, but in others, more risk can be accepted because of the materiality of the model. The FI should use the information provided by the vendor to as a basis for their own reasoning during the model validation stage.
- Accountability for the deployed system → Accountability presupposes transparency, and while there is a general agreement that the FI is accountable for the deployed technology, this is conditional, depending on the full clarity provided by the vendor around the methodology used and the increased focus on explainability and interpretability. This



includes overall end-to-end adopted modelling techniques, any assumptions or approximations that were made, and details of the processing stages – with a focus on the configuration that has impact on the overall model output. While the vendor might provide some default values used in the configuration, and the reasoning for such values, it is critical to provide clarity and on how those parameters can be tuned. The vendor should also support the FI by describing, for each parameter, impact on the overall model output, potential reasoning for the acceptance of impact, and possible methods to mitigate the risk. Together with the possible mitigations related to the general limitations of the system, this would provide input for the FI's work on designing appropriate control frameworks for the technology.

- Traceability for the final configuration of the deployed system → Both vendors and FIs need to clearly document the decisions made in every stage of the lifecycle, but especially those made during the deployment. Vendors need to provide justification for any default values, including the guidance, expertise, or data analysis behind it, and FIs have to document the process of accepting or tuning these parameters to the production values, including the roles that made the decisions, when those decisions were made, and the input and supporting evidence that formed the basis of the decisions. This information needs to be compiled in comprehensive reports that will support the model validation processes. Although vendors can provide guidance and expertise based on their understanding of the technology and their experience delivering the technology, FIs are ultimately responsible for the configuration of their solutions, and any vendor inputs should be clearly documented as such. It is critical that all parameters are therefore appropriately understood, analysed, and tuned as required by the FI (or a third party engaged for the task).
- Data governance for all types of datasets used → While the FI is responsible for the procedures in place regarding production data, they might have to rely on the vendor's procedure regarding training and evaluation data for the vendor products. Vendors should be able to present information regarding the training and evaluation datasets and provide documentation related to data pre-processing, data quality, and relevance analysis. If required by the FI, the vendor should be able to provide access to the datasets when possible, considering data privacy and licensing restrictions. For any dataset used separately from the actual implementation, the vendor must support the FI with information required by their data governance procedures to confirm that the dataset went through a governance process that, while not identical to that of the FI, respects the same principles.
- Testing & evaluation of the deployed system → Similar to data governance, the vendor should provide the FI with the testing and evaluation protocols, an overview of what has been measured in the evaluation, and the results from product development. By documenting the diagnostic, performance, and any other tests performed across development data and validation data – including an overview of which metrics are tracked during model monitoring – the vendor provides the FI with the assurance required for model validation. The use of parallel runs to operate a new or challenger model alongside the existing (champion) model might introduce significant challenges when it



comes to advanced technologies as a result of potential differences of transparency or explainability gaps and inability to completely align the outputs; FIs can rather focus on the performance of each system to achieve the desired objectives while factoring in the impact of a new system on the end-users' performance.

The MRM roles involved in the deployment need to be closely embedded in the team not only to make sure that all information required for model documentation is collected, but also that any modelling decisions made during the deployment (e.g., tuned parameters) are documented with a clear supporting rationale and data. This is also critical to ensure that the appropriate controls are designed to complement the technology as it is actually deployed, and not just expected to be deployed.

IMPLEMENTATION / DEPLOYMENT - QUESTIONNAIRE

13. What is the design, theory, and logic of the model, including the key steps of the modelling process and the algorithms used?
14. What are the modelling decisions that drove model design work (including product development decisions and configuration decisions)?
15. How has the final configuration of the deployed system been achieved?
16. What datasets have been used (or are planned to be used), and what data governance procedures were applied?
17. What are the testing and evaluation protocols applied, including the results pre-production? Which ones will be used for monitoring?

Production

Once the technology is running in the production environment, significant attention must still be paid as part of business as usual processes. FIs should periodically evaluate whether the project continues to meet requirements and to check whether the area of application has been changed.

It's necessary to embed the following topics in the FI's regular activities in order to maintain the right level of governance throughout the regular technology run:

- Performance monitoring of the system → FIs generally have clear internal procedures that indicate the requirements for model monitoring, like cadence and triggering metrics. These should be aligned to the release cadence of the vendor and the planned upgrade approach, since upgrades might come with modifications to the model reflected by the technology. When there is strong collaboration from a model governance perspective, the FI could request that the vendor highlight in the release notes if the changes are impacting modelling and output in a way that could trigger a model review. The FI is ultimately responsible for the actual decision to trigger an event-based model review, but it can engage the support of the vendor as it does in the other lifecycle stages.



- Total cost of ownership over the complete lifecycle → Unlike traditional IT projects, where costs are high during deployment and drop significantly after implementation, AI projects generally require more model maintenance, and present costs that scale with usage. FIs need to make sure the vendor cost structure is transparent and includes also the costs indirectly related to the vendor's product incurred by its interfaces with other products (e.g., processing units, databases). They also need to account for the model monitoring approach recommended or applied by the vendor (since retraining and upgrades come with significant cost), and for the release cadence that brings potential new value at the clear cost of a technical upgrade.
- Achieving and measuring of desired outcomes → While success outcomes and metrics should have been discussed and agreed upon in the earlier phases of the deployment, their achievement has many dependencies, including the effective operationalisation of the technology by the FI, as well as a clear articulation of the state of play prior to technology deployment. As such, even with sufficient preparation and effective operating model changes, there can still be instances where vendors can help the FIs improve their success outcomes and ensure the deployed technology realises its maximum value. Tracking the right metrics should be accompanied by performing an educated deep analysis of root causes and implications to drive a clear view of supporting actions. FIs can benefit from vendor expertise in this space until they reach self-sufficiency.

An ongoing collaboration between FI and vendor on model risk management activities can provide value for both parties: FIs can better align their internal monitoring processes with the vendor's release cadence, and vendors can use the feedback to improve the AI system in case of performance drift on the FI side. At the same time, the actual effectiveness of the risk controls deployed by the FI can be used as user feedback for product development. As long as proper privacy is built into the collaboration, we recommend regular check-ins between the model governance representatives of the vendor and the FI.

PRODUCTION - QUESTIONNAIRE

18. How are you monitoring the performance of the model, and when are you expecting a model review?
19. What are the cost elements of the total cost of ownership, and how are you tracking them over time?
20. How are you measuring the metrics linked to the desired outcomes of the technology?



CONCLUSION / CALL TO ACTION

While FIs are ultimately accountable for the effective risk management of the technologies they deploy, the growing complexity of these solutions and increasing reliance on third-party providers highlight the need for stronger alignment and more structured collaboration between FIs and technology vendors in managing technology risk. As highlighted in the International Chamber of Commerce (ICC) Jul 2025 policy paper on AI governance and standards¹³, market-driven standards can play an important role in providing practical solutions and guidance on how to comply with laws, policies and regulation while reducing duplication and improving regulatory coherence.

This paper proposes a framework to help both FIs and vendors to establish clearer operational risk standards across the technology lifecycle, bridging the gap between deep technical solution expertise and robust MRM processes.

To drive meaningful progress, we recommend the following actions:

- **The FI industry should seek to standardise MRM requirements** to reduce friction and enable more consistent vendor engagement. Buy-in and endorsement from relevant industry organisations could materially help drive and accelerate this process.
- **Vendors should align with FI requirements** by developing industry-standard documentation and support practices based on the principles outlined in this framework. Refer to *Appendix 1 – Documentation Standards* for a baseline of a Vendor Model Card.
- Following the establishment and testing of FI and vendor standards, **vendors** are also encouraged to go a step further and **integrate MRM capabilities directly into their solutions**, truly embedding MRM as part of product design and enabling future technology innovation in the MRM space (e.g., MRM agents).

By adopting a shared set of best practices, FIs and vendors can strengthen technology risk management, improve efficiency, and enable more scalable and resilient advanced technology adoption across the industry.

¹³ [ICC \(2025\), ICC Policy paper on AI governance and standards](#)



Appendix

Appendix 1 – Documentation Standards

Vendor Model Cards to align to FI MRM needs

For MRM purposes we provide an example of standard documentation:

- **Data Card** → Data-focused documentation applied to datasets used by the vendor in model development
- **ML Model Card** → Models-and-methods-focused documentation applied to machine learning (ML) models and methods (template provided below)
- **Component Card** → Technology-focused documentation applied to decision-based rules or algorithms (i.e., technology products not using ML)
- **Solution Card** → Solution-focused documentation applied to the needs and configurations that are specific to a use case (e.g., the application of a certain vendor product for financial crime compliance)
- **System Card** → Systems-focused documentation applied to a group of AI and non-AI technologies, including non-ML models, which work together to accomplish specific tasks (e.g., the deployment of a certain vendor product by a user, including all local configurations)

ML Model Card Example

Details / Summary

This card has been prepared leveraging existing technology industry standards¹⁴ but has been adapted to closer align to FI requirements.

This section of the ML model card should serve to answer basic questions regarding the model version, type, and other details.

Section	Description
Name & description	Provide the ML model name and a 1–2 sentence summary of the nature of ML model.
Version & dates	Indicate the version of the ML model and how it differs from previous versions. This is to enable all stakeholders to track whether the ML model is the latest version, associate known bugs to the correct ML model versions, and aid in ML model comparisons. Indicate when the ML model was developed and when it was released. Include details on the update cadence.

¹⁴ [THE LANDSCAPE OF ML DOCUMENTATION TOOLS](#)



Purpose & Usage	Describe the purpose and products for which the ML model is designed, including actual or expected usage. Include: • summary of ML model background (including problem statement) and summary of the product / portfolio / population to which the model will be applied. • the business process in which the ML model is embedded, and explain how it integrates into this process. • the key objectives of the ML model (e.g., automate process to achieve operational cost savings, reduce manual efforts, or enhance customer and user experience, etc.) and the success criteria used to measure the key objective(s). Include any restrictions on use or other controls (e.g., more frequent monitoring and appropriate benchmarking) and out-of-scope uses. If possible, for out-of-scope uses, indicate a related or similar component or ML model that was designed to better meet that particular need.
Techniques	Provide a high-level description of the type of technology or approach used. Where applicable, include comparison with alternative theories and approaches.
Input(s)	Indicate the type and source of inputs used by the ML model and its underlying elements (which may include other components or ML models).
Output(s)	Describe ML model outputs and their intended use.
Scope	Describe the applicable scope for the ML model, including relevant data, geographic scope, and population. Include any other boundaries within which ML model performance is expected to be acceptable.
Deployment type	Indicate if this is a standalone ML model or intended to be used as part of a system with other components / ML models. Include links where necessary. Explain how the ML model can be used without fine-tuning, post-processing, or plugging into a pipeline. Explain how this ML model can be used when fine-tuned for a task or when plugged into a larger ecosystem or app. Include upstream dependencies (If the ML model requires specific inputs, where should they come from? Are there any specific preprocessing steps that should be applied?) or downstream dependencies (If the ML model's outputs can be fed into another system, where should they go? Are there any specific post-processing steps that should be applied?).
Deliverable type	Describe the nature of the deployable ML model (e.g., a set of predictions, a pretrained scoring function with parameters, a modelling pipeline, or methodology recommendations).
License & proprietary information	List the license requirements applicable to the ML model (e.g., links to external data). Indicate if there are any license constraints for use of the ML model (especially data licensing constraints), and if there is any information, including but not limited to code, documentation, and parameters, which cannot be freely shared with appropriately licensed customers.
Alternative approaches	List any existing alternative methods used to achieve the ML model objective. Include performance measures for such rules when available.



Origin	Indicate if the ML model was developed in-house or by a third party. Provide reasoning why a third-party ML model was selected. Include internal IP classification.
--------	---

Theory and Technical Specifications

This section includes details about the ML model objective and architecture, and the compute infrastructure.

Section	Description
Methods / Algorithms included	Describe all algorithms, including those used for feature generation or cleansing. Include links to definitions of any non-typical algorithms. Evaluate the applicability of selected algorithms against the business objective and ML model use. Provide reasoning on selection that consider the accuracy vs. explainability trade-off.
Modelling pipeline	Describe the design, theory and logic of the ML model. Describe the key steps from the modelling process. List all stages applied to the raw source data, including those to create a modelling data set and train the model, including validation. Include link to code.
Modelling assumptions	List all assumptions behind the ML model, including both mathematical assumptions behind the selected algorithms and other assumptions made during the modelling process.
Technology	Describe the hardware and software used for training the ML model.
Compute requirements	Describe the following compute requirements, where applicable: Number of Chips, Training Time (days), Total Computation (floating point operations), Measured Performance (total floating point operations per second), and Energy Consumption (megawatt-hours) (especially for Large Models).

Data Overview

This section provides information regarding the training and evaluation data. Links to documentation related to data pre-processing or additional filtering may go here.

Ideally, the model card would contain as much information about the training data as the evaluation data. However, there might be cases where it is not feasible to provide that level of detailed information about the training data. For example, the data may be proprietary or require a non-disclosure agreement. In these cases, we advocate for basic details about the distributions over groups in the data, as well as any other details that could inform stakeholders on the kinds of biases the model may have encoded.



Section	Description
Data sources	Describe the dataset used in model training and evaluation; indicate separately the exact use per each dataset (e.g., training, initial testing, additional testing). Consider dimensions like: Dataset Size, Number of Instances, Number of Fields, Labelled Classes, Number of Labels, Average Labels per Instance, Missing Labels. • Type (trusted / tactical, internal / external) • Time period • Description • Applicability to modelling (i.e., how the data is used within the model) • Data assumptions / limitations • Impact on model development • Critical data elements (CDEs) • Data exclusions • Jurisdictions covered In case proxy data is used, include the description of proxy data and justification of relevance for modelling objective. Include link to retained training dataset.
Dataset maintenance & versions	Indicate if the training data is static or updated/expanded. If so, indicate the frequency with which this data is updated.
Data quality and relevance analysis	Indicate what data quality checks were applied, with summary results. Ensure coverage for the applicable dimensions.
Demographic groups	Indicate if the data contains any labelled groups, or attributes that suggest demographic group membership. Describe any demographic groups considered when assessing distributions in the data. If there are groups that may be present, but are not labelled in the training data, note this in the Risks & Limitations section.
Representativeness analysis	List all representativeness checks, with summary results. Describe the data profiles included in the modelling data – counts of records selected for training/validation by region, by date, and by any other relevant metric for the model type (e.g., industry for a business model).



Methodology Overview

This section includes the overall end-to-end modelling techniques adopted and assumptions or approximations made, and details of the processing components used in the ML model implementation. It covers all stages from the modelling process in detail.

Section	Description
Data pre-processing	Describe any augmentation methods used during pre-processing to attain the requisite format. List the criteria that data points must satisfy to be included in the training set, if applicable. Describe pre-processes in place (e.g., expanding contractions and lower casing, removal of numbers and special characters and punctuation, removal of stop words, stemming / lemmatisation, removal of frequent and rare words, removal or conversion of emoticons and emojis to words, removal of low-quality data, etc.), the reason for each of them and the data sources to which any technique is applied. If filters are used, counts should be included before and after filtering.
Dependent / Target variable(s)	Describe the target variables for the ML model (if it's a supervised model) and the process used for obtaining labels. Describe the quality of the labelled data and propose controls to maintain the accuracy of the process, if applicable (e.g., human labelling).
Feature engineering	Describe pre-processing of raw data into interpretable features (rather than solely relying on the ML algorithm) for quantitative (non-vectorised) input variables. In case the features are calculated at different levels, include the feature aggregation method.
Feature selection	Describe the feature analysis and selection process, including reasoning.
Training hyperparameters	Include a summary of hyperparameters tuned, range of hyperparameter values tested, optimisation technique (e.g., grid search, random search, or Bayesian optimisation), metrics and criteria used to select the "optimal", along with a summary of results of tested iterations.
Evaluation functions	Describe the selected objective function and evaluation metrics used for ML model training and evaluation.
Final specifications	Provide codes, hyperparameter value / configuration settings, and random seed values used in ML model development, to ensure reproducibility.
Alternative algorithms / methods considered	List all algorithms evaluated, with reasons for inclusion or exclusion.



Testing & Evaluation

This section describes the testing and evaluation protocols, as well as what is being measured in the evaluation, and provides the results. Evaluation is ideally constructed with factors (such as domain and demographic subgroup) and metrics (such as accuracy), which are prioritised considering foreseeable error contexts and groups. Target fairness metrics should be decided based on which errors are more likely to be problematic in light of the model use.

Section	Description
Testing approach	Provide a test approach which includes the diagnostic test, performance test, and any other tests performed. The type of tests depends on the modelling approach and intended uses. The performance test should be performed across development data (training data) and validation data (test and out-of-time data). Include the approach used to split the different datasets. The section should also indicate which tests are to be performed as part of monitoring. List the generic tuning & monitoring processes, if any, and indicate the out-of-the-box tools that support them.
Soundness / diagnostic results	Include results related to the ML model soundness / diagnostic tests with supporting calculation files and conclusions. Include any analysis performed on output that is not labelled data.
Performance evaluation results	Include results related to the ML model performance tests, as well as supporting calculation files and conclusions. Include link to retained scored ML model validation dataset, allowing recreation or extension of validation test results.
Subgroup performance evaluation results	Document your disaggregated evaluation. Duplicate the following for each subgroup evaluated: • subgroup evaluated: indicate the evaluated subgroup. • evaluation process and data: describe any notable factors in your process for disaggregated or sliced evaluation of model performance; include any assumptions made when disaggregating the data. • evaluation results: indicate any known and preventable failures about the model. Include Group versus Individual fairness measurements (i.e., disparate impact) and consider other fairness metrics, such as equalised odds or equality of opportunity, that may be more appropriate for some applications/domains than disparate impact.



Risks and Limitations

This section identifies foreseeable harms, misunderstandings, and technical and sociotechnical limitations. It also provides information on warnings and potential mitigations. Bias, risks, and limitations can sometimes be inseparable, or refer to the same issues.

Section	Description
Limitations	List self-identified limitations, including proposals to address or remediate the concerns (e.g., through monitoring, targeting only eligible population in the model deployment, ensuring accuracy of labelled dataset for future validation workflows, etc.).
Sensitive use	Indicate if there are use cases where deployment of this ML model would be considered sensitive. Sensitive uses directly impact how a party is treated, such as raising an alert or closing an account. Examples of non-sensitive uses would include how data is processed, such as classifying a name as a business, where model bias could not directly cause adverse impact. Map the potential and recommended use cases to the use cases included in the EU AI Act (or other applicable laws or regulatory requirements).
Sensitive data	Indicate if any of the data inputs are related to protected characteristics or are potentially likely to act as a proxy for such characteristics. Include mitigation or justification in cases where the data input is used for modelling purposes.
Risks	Indicate the risks, including planned or possible mitigations.

Deployment

This section includes details about ML model use.

Section	Description
Technical dependencies	List the detailed technical dependencies required to deploy the ML model.
Deployment approach	Provide a description of how the ML model should be deployed.
Smoke test	Share an environment test that can be used to test environments prior to ML model deployment.
Unit test	Provide code and test data to validate that the ML model is performing identically on a new environment as compared to the original.



Appendix 2 – Technology Impact Assessment Questionnaire

Stage	Questionnaire
Problem / Opportunity Identification & Solution Selection	1. What is the intended purpose and result of the technology? 2. What is the intended solution to address stated purpose and result? 3. What are the prerequisites for the intended solution? 4. What is the overall maturity of the features included in the intended solution? 5. Are there any alternative solutions to address the described purpose and result? 6. How will the intended solution be deployed? 7. What are the licenses required for the use of the solution, and what are the licenses linked to the assets used in the development? 8. How is compliance with applicable laws and regulation being achieved?
Mobilisation	9. What is the risk tier allocated to the model, and what are the associated considerations? 10. What are the potential risks and limitations associated with the targeted technology? 11. What are the specialised skills required, and how are they sourced? If an external party is responsible for the deployment, what contractual agreements are in place? 12. What needs to change in the operating model to drive the targeted value outcomes?
Implementation / Deployment	13. What is the design, theory, and logic of the model, including the key steps of the modelling process and the algorithms used? 14. What are the modelling decisions that drove model design work (including product development decisions and configuration decisions)? 15. How has the final configuration of the deployed system been achieved? 16. What datasets have been used (or are planned to be used), and what data governance procedures were applied? 17. What are the testing and evaluation protocols applied, including the results pre-production? Which ones will be used for monitoring?
Production	18. How are you monitoring the performance of the model, and when are you expecting a model review? 19. What are the cost elements of the total cost of ownership, and how are you tracking them over time? 20. How are you measuring the metrics linked to the desired outcomes of the technology?



Appendix 3 – MRM / AI Vendor Requirements Overview

List and extracts of select global MRM / AI guidelines and requirements as pertaining to vendors / third parties

Geo	Regulation / guidance	Primary requirements related to vendors
US	SR 11-7: Guidance on Model Risk Management, Federal Reserve, 2011	Scope: banks / regulated institutions – broad model risk management. Extracts: <i>Validation of Vendor and Other Third-Party Products</i> The widespread use of vendor and other third-party products – including data, parameter values, and complete models – poses unique challenges for validation and other model risk management activities because the modeling expertise is external to the user and because some components are considered proprietary. Vendor products should nevertheless be incorporated into a bank's broader model risk management framework following the same principles as applied to in-house models, although the process may be somewhat modified. As a first step, banks should ensure that there are appropriate processes in place for selecting vendor models. Banks should require the vendor to provide developmental evidence explaining the product components, design, and intended use, to determine whether the model is appropriate for the bank's products, exposures, and risks. Vendors should provide appropriate testing results that show their product works as expected. They should also clearly indicate the model's limitations and assumptions and where the product's use may be problematic. Banks should expect vendors to conduct ongoing performance monitoring and outcomes analysis, with disclosure to their clients, and to make appropriate modifications and updates over time. Banks are expected to validate their own use of vendor products. External models may not allow full access to computer coding and implementation details, so the bank may have to rely more on sensitivity analysis and benchmarking. Vendor models are often designed to provide a range of capabilities and so may need to be customized by a bank for its particular circumstances. A bank's customization choices should be documented and justified as part of validation. If vendors provide input data or assumptions, or use them to build models, their relevance for the bank's situation should be investigated. Banks should obtain information regarding the data used to develop the model and assess the extent to which that data is representative of the bank's situation. The bank also should conduct ongoing monitoring and outcomes analysis of vendor model performance using the bank's own outcomes. Systematic procedures for validation help the bank to understand the vendor product and its capabilities, applicability, and limitations. Such detailed knowledge is necessary for basic controls of bank operations. It is also very important for the bank to have as much knowledge in-house as possible, in case the vendor or the bank terminates the contract for any reason, or if the vendor is no longer in business. Banks should have contingency plans for instances when the vendor model is no longer available or cannot be supported by the vendor.
US	Artificial Intelligence Risk Management Framework (AI RMF 1.0)	Scope: AI actors (broad range of stakeholders including individuals and organisations) - AI systems defined as: “an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy (Adapted from: OECD Recommendation on AI:2019; ISO/IEC 22989:2022)”.



		Given the scope, the AI RMF applies to all firms developing and deploying AI and is not limited to banks / regulated institutions. Nevertheless, the RMF specifically addresses reliance on vendors / third-parties and this is included below. Extracts: <i>Risks related to third-party software, hardware, and data:</i> Third-party data or systems can accelerate research and development and facilitate technology transition. They also may complicate risk measurement. Risk can emerge both from third-party data, software or hardware itself and how it is used. Risk metrics or methodologies used by the organization developing the AI system may not align with the risk metrics or methodologies uses by the organization deploying or operating the system. Also, the organization developing the AI system may not be transparent about the risk metrics or methodologies it used. Risk measurement and management can be complicated by how customers use or integrate third-party data or systems into AI products or services, particularly without sufficient internal governance structures and technical safeguards. Regardless, all parties and AI actors should manage risk in the AI systems they develop, deploy, or use as standalone or integrated components. <i>Specific reference under AI RMF Core:</i> GOVERN 6: Policies and procedures are in place to address AI risks and benefits arising from third-party software and data and other supply chain issues. GOVERN 6.1: Policies and procedures are in place that address AI risks associated with third-party entities, including risks of infringement of a third-party's intellectual property or other rights. GOVERN 6.2: Contingency processes are in place to handle failures or incidents in third-party data or AI systems deemed to be high-risk. MAP 4: Risks and benefits are mapped for all components of the AI system including third-party software and data. MAP 4.1: Approaches for mapping AI technology and legal risks of its components – including the use of third-party data or software – are in place, followed, and documented, as are risks of infringement of a third party's intellectual property or other rights. MAP 4.2: Internal risk controls for components of the AI system, including third-party AI technologies, are identified and documented. MANAGE 3: AI risks and benefits from third-party entities are managed. MANAGE 3.1: AI risks and benefits from third-party resources are regularly monitored, and risk controls are applied and documented. MANAGE 3.2: Pre-trained models which are used for development are monitored as part of AI system regular monitoring and maintenance.
UK	PRA supervisory statement SS1/23	Scope: banks / regulated institutions – broad model risk management (covering all models developed in-house or externally, including vendor models, and models used for financial reporting purposes). Extracts: <i>Principle 2.6 Use of externally developed models, third-party vendor products</i>



		a) In line with PRA SS2/21 – Outsourcing and third party risk management boards and senior management are ultimately responsible for the management of model risk, even when they enter into an outsourcing or third-party arrangement. b) Regarding third-party vendor models, firms should: (i) satisfy themselves that the vendor models have been validated to the same standards as their own internal MRM expectations; (ii) verify the relevance of vendor supplied data and their assumptions; and (iii) validate their own use of vendor products and conduct ongoing monitoring and outcomes analysis of vendor model performance using their own outcomes. <i>Principle 3.5 Model development documentation</i> Firms should ensure the level of detail in the documentation of third-party vendor models is sufficient to validate the firm's use of the model.
EU	ECB guide to internal models	Scope: banks / regulated institutions – internal models that are subject to supervisory approval for the calculation of own funds requirements for credit, market and counterparty credit risk. Due to the more targeted scope of the guide, no extracts are provided, however <i>Chapter 8 Third-party involvement</i> covers requirements pertaining to third-parties and outsourcing.
EU	The Act Texts EU Artificial Intelligence Act	Scope: AI providers / deployers / importers and distributors / product manufacturers / authorised representatives / AI users (natural or legal persons that deploy an AI system in a professional capacity) – AI systems defined as “<i>a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments</i>”. The Act does not differentiate between the type of AI providers and therefore would apply to any natural or legal person engaged in the regulated activities. However, at present most AI models used in banks are likely to fall out of scope of the EU AI act with the possible exception of a) HR systems and b) retail / wealth management / private banking credit / pricing models. Refer to the list of high-risk systems: Section 1: Classification of AI Systems as High-Risk EU Artificial Intelligence Act Annex III: High-Risk AI Systems Referred to in Article 6(2) EU Artificial Intelligence Act
AE	CBUAE Model Management Standards	Scope: Banks – Models (all models used to support decision-making). Extracts: <i>4.7 Third party provider</i> 4.7.1 Institutions must remain the owners of their models at all times, under all circumstances. They must remain accountable for all modelling choices, even in the case of support from a third party consultant for any of the steps in the life-cycle. 4.7.2 If modelling support is provided by a third party, institutions must take the necessary steps to transfer knowledge from that third party to internal employees within a given time frame. This requirement applies to any of the steps of the model life-cycle. 4.7.3 Third party providers may offer a range of modelling contributions covering, amongst others, methodological support, system infrastructure, validation services and ready-made



		calibrations based on external data. Institutions must take the necessary action to fully understand the contributions provided by third parties. This requirement applies to all models and to all risk types. 4.7.4 In the case of methodological support, whilst institutions must operate within the constraints of the acquired model, they must demonstrate that the method is adequate to their portfolios. If a methodology acquired from a third party is not fully understood by the institution, then it must not be considered fit for purpose. If a third party provides a methodology to an institution, any subsequent validation exercise must be performed by an internal or external party independent from the original provider. 4.7.5 If a third party provides a ready-made calibrated model based on external data, such a solution must be justified, based on the following specific circumstances: (i) For portfolios and metrics for which an institution is not able to collect sufficient internal data, then externally calibrated models are acceptable. For instance, this applies in the case of low default portfolios or small portfolios for which data collection may not lead to statistically representative samples. (ii) For portfolios and metrics for which an institution is in a position to collect internal data, then externally calibrated models must not be used. Externally calibrated models are acceptable, only temporarily over the short term until sufficient data is collected. In this case, immediately after the model implementation, institutions must take the necessary actions to (i) collect historical internal data from internal systems and (ii) collect future internal data in order to develop a model internally.
AE	UAE AI Ethics Principles & Guidelines	Scope: AI design / development in private and public sectors – AI systems used for significant decisions (decisions which have the potential for significant impact either on individuals or on society as a whole). AI is defined as <i>‘the capability of a functional unit to perform functions that are generally associated with human intelligence such as reasoning, learning and self-improvement’</i>. AI system is defined as <i>‘a product, service, process or decision-making methodology whose operation or outcome is materially influenced by artificially intelligent functional units’</i>. Extracts: N/A
HK	HKMA High-level Principles on AI	Scope: Banks – AI (undefined) Extracts: <i>Implementing effective management oversight of third-party vendors</i> – Where banks rely on third-party vendors to develop AI applications, they should perform proper due diligence on these vendors having regard to the applicable principles set out in this letter. They should also implement effective vendor management controls including periodic reviews of the services provided to manage the associated risks.
SG	MAS AI Model Risk Management	Scope: Banks / regulated institutions – AI (including GenAI), defined as <i>“a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment (based on the Organisation for Economic Cooperation and Development’s definition of AI). Such a definition would include Generative AI. An AI or Generative AI system can be based on one or multiple AI or Generative AI models and may also involve other machine-based components.”</i> Extracts:



		<i>7.2 Third-Party AI</i> <i>Overview</i> Existing third-party risk management standards and processes continue to play an important role in banks' efforts to mitigate risks associated with third-party AI. As far as practicable, most banks also extended controls for internally developed AI to third-party AI. When considering the use of third-party AI, banks would weigh the potential benefits against the risks of using third-party AI. To address the additional risks arising from third-party AI, banks were exploring areas such as: • conducting compensatory testing; • enhancing contingency planning; • updating legal agreements; and • investing in training and other awareness efforts. 7.2.1 The use of third-party AI is increasingly common among banks, particularly in the context of Generative AI where most banks utilise Generative AI models that were pre-trained by an external party. However, the use of such third-party AI and Generative AI presents additional risks, such as unknown biases from pre-training data, data protection concerns, as well as concentration risks due to increased interdependencies, e.g., from multiple FIs or even third-party providers relying on common underlying Generative AI models. The lack of transparency is often cited as a key challenge in managing such third-party risks. Third-party AI providers may be reluctant to disclose proprietary information about their training data or algorithms, hindering banks' efforts in risk assessment and ongoing monitoring. 7.2.2 To mitigate these additional risks, banks were exploring various approaches, such as: a) Compensatory testing - conducting rigorous testing of third-party AI models using various datasets and scenarios to verify the model's robustness and stability in the bank's context, and to detect potential biases. b) Contingency planning - developing robust contingency plans to address potential failures, unexpected behaviour of third-party AI, or discontinuing of support by vendors. This can include having backup systems or manual processes in place to ensure business continuity. c) Legal agreements- updating contracts with third-party AI providers to include clauses such as those pertaining to performance guarantees, data protection, the right to audit, and notification when AI is introduced (or not incorporating AI without the bank's agreement) in existing third-party providers' solutions. Such clauses could facilitate clearer expectations and responsibilities. Awareness efforts – investing in training of staff on AI literacy and risk awareness to improve understanding and mitigation of risks; conducting surveys with third-party providers to gather more information about whether AI is being used in their products or services, and third-party providers' practices, including their AI development and risk management processes.
AU	Department of Industry, Science and Resources Voluntary AI Safety Standard	Scope: AI deployers (primarily) / AI developers – AI systems defined as “A machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment (OECD).” Extracts: 1.1.3 Clearly communicate the leadership commitment to, and accountability for, safe and



	responsible development and use of AI across the organisation. This includes the staff (including contractors and third-party providers) who you have made accountable for AI systems. 1.2.7 Create and document a process for deploying AI systems that supports mapping from business targets to system performance, with suggested metrics for internal and third-party developed systems. 1.3.5 Where applicable, evaluate the training needs for staff who deal with third-party AI systems that are being developed, procured or used. Provide the appropriate training to address skill gaps. 2.1.3. Create and document a suitable impact assessment, risk assessment and treatment approach to AI system deployment and use. This should cover both internal and third-party developed AI systems, with awareness of the specific characteristics and amplified risks of AI systems. Include criteria for reassessment over the lifecycle of an AI system. 5.1.6. Assign accountability for oversight of third-party development and use of AI systems and components to appropriately skilled and empowered people in the organisation. 6.2.1. Evaluate the level of transparency that each AI system needs – including third-party-provided systems – dependent on the use case and external stakeholder expectations. Consider potential conflicts, such as privacy, intellectual property, AI systems presenting as a person, hallucinations or potential for misinformation. 6.2.5. Where expected by stakeholders, implement approaches to communicate relevant information about AI-generated content to end users. Require associated third-party developers to do the same, with options such as labelling and watermarking. Evolve these approaches as new solutions become available. 9.2.7. Ensure documentation related to each AI system is recorded in the inventory at a sufficient and consistent level of detail to inform the accountable and responsible parties and any third-party stakeholders. This will enable completion of future conformity assessments to demonstrate compliance with mandated guardrails. d) <i>Procurement guidance for guardrail 9: Work with your supplier to understand and document the expected use, capabilities and limitations of the AI system or component. This should include technical details of the system and the data used in relation to the AI system (including the use of third-party data). Integrate expectations into contract, including ongoing scheduled reviews.</i>
--	---