OAK & HILL, P.C.
ATTORNEYS AT LAW

AI LEGAL INFRASTRUCTURE

The ARIA Playbook

A Legal Risk Guide for AI-Integrated Companies

What you will find inside:

- A structured review of the five legal risk layers in every AI-integrated business
- Plain-English explanations of how exposure accumulates across data, contracts, and deployment
- Self-assessment checklists for each risk domain
- A framework for prioritizing remediation before the next funding round, acquisition, or regulatory inquiry
- A clear next step when you are ready to move from awareness to structure

Contents

Introduction	How to Use This Guide	3
Overview	The ARIA Framework at a Glance	5
Layer A	Data Architecture — Who Owns Your Training Data?	7
Layer R	Risk Allocation — Where Does Liability Land?	10
Layer I	IP and Output Ownership — Who Owns What the AI Creates?	13
Layer A_2	Accountability and Governance — Can You Prove You Were Careful?	16
Layer 2	Regulatory Architecture — What Frameworks Apply to You?	19
Closing	Self-Assessment Summary and Next Steps	22

Introduction

How to Use This Guide

This guide is not a compliance checklist. It is a diagnostic tool.

Most companies integrating artificial intelligence into their products, operations, or workflows are doing so faster than their legal structures can keep pace. That gap is not hypothetical risk. It is documented exposure accumulating in vendor agreements that were never negotiated, in employment arrangements that never addressed AI-generated work product, in data pipelines that operate on assumptions rather than rights, and in deployment architectures that allocate liability to no one.

The ARIA Framework—developed by Oak & Hill, P.C. as the organizing structure for our AI legal practice—was designed to bring that gap into focus before it becomes a problem in a funding round, a due diligence process, a regulatory inquiry, or litigation.

ARIA stands for five sequential layers of legal risk that every AI-integrated company carries, whether or not it has addressed them:

A DATA	Data Architecture — the legal rights, restrictions, and obligations governing the data your AI uses
R RISK	Risk Allocation — how liability is distributed across your AI value chain
I IP	IP and Output Ownership — who legally owns what your AI creates
A_2 GOV.	Accountability and Governance — whether your organization can demonstrate responsible AI use
2 REG.	Regulatory Architecture — which legal frameworks govern your AI operations, now and prospectively

How This Guide Is Structured

Each of the five chapters that follow covers one ARIA layer. Each chapter contains three components: a plain-English explanation of the risk domain, a description of how exposure commonly accumulates in practice, and a self-assessment checklist that allows you to evaluate your current position.

At the end of the guide, a consolidated self-assessment summary gives you a prioritized view of where your company stands across all five layers.

This guide is written to be read by founders, operators, and senior leaders—not just legal teams. The risk areas covered here are commercial and operational in nature. The decisions required to address them are not exclusively legal decisions. They require judgment about business structure, vendor relationships, data strategy, and governance that spans the organization.

A Note on Scope

This guide addresses U.S.-based companies, with relevant notes on cross-border exposure where the risk is material. It does not constitute legal advice, and it does not create an attorney-client relationship between Oak & Hill, P.C. and the reader. Its purpose is to surface the questions every AI-integrated company should be asking—and to provide a structured framework for addressing them with competent counsel.

If you find, after completing this guide, that your company has unaddressed exposure in multiple layers, the appropriate next step is a formal AI Legal Architecture Review—a structured engagement through which Oak & Hill delivers a documented, prioritized analysis of your specific situation.

The ARIA Framework at a Glance

Before examining each layer in depth, it is useful to understand the relationship between them. The ARIA layers are not independent checklists. They are interconnected risk domains where decisions made in one layer create obligations and vulnerabilities in others.

A company that has excellent data provenance documentation (Layer A) but no allocation of output liability in its commercial agreements (Layer R) has closed one exposure while leaving another wide open. A company with comprehensive internal governance documentation (Layer A2) but no analysis of the EU AI Act's applicability to its product (Layer 2) is well-organized for a problem it may not face, while unprepared for one it will.

The framework is intended to be applied as a whole. Partial application produces partial protection.

Where Most Companies Are Today

Risk Area	Exposure Without Structure	Priority
Data Rights & Provenance	Using training or fine-tuning data without confirmed rights; signed platform agreements that transfer data to vendors	Critical
Vendor Liability Allocation	No negotiated AI-specific provisions in SaaS, API, or platform contracts; vendor terms disclaim all liability for outputs	Critical
Output IP Ownership	AI-generated content delivered to customers with undefined ownership; no work-for-hire or assignment provisions in place	High
Internal AI Governance	No formal AI use policy; employees using AI tools without documented oversight; no incident response protocol	High
Regulatory Exposure	No analysis of CCPA/GDPR implications for AI data pipelines; EU AI Act applicability not evaluated; sector-specific rules ignored	Medium
Employment & IP Assignment	Offer letters and IP agreements predate AI; contractor work product involving AI not covered by assignment clauses	Medium
Cross-Border Data Transfer	AI training data or inference processed across jurisdictions without transfer mechanisms in place	Monitor

The table above reflects patterns observed across early- and growth-stage companies that have adopted AI tooling without corresponding legal review. It is not a worst-case scenario. It is a common baseline.

The chapters that follow will allow you to locate your company precisely within each of these domains and identify the priority actions required to move from exposure to structure.

LAYER A Data Architecture

Who Owns Your Training Data—and What Did You Promise When You Got It?

Every AI model is a function of the data used to train, fine-tune, or instruct it. The legal questions surrounding that data are foundational: who owns it, under what terms was it obtained, what restrictions attach to its use, and what has the company represented—expressly or implicitly—to the people whose information it contains?

For most companies, the honest answer to at least one of those questions is: we are not certain. That uncertainty is not a minor operational gap. It is the kind of issue that surfaces at the worst possible moment—during diligence for a funding round, in response to a regulatory inquiry, or when a vendor or data partner disputes the scope of their license.

How Data Risk Accumulates

Third-Party Training Data

Companies that use publicly scraped data, licensed datasets, or data obtained through APIs to train or fine-tune models frequently do not have confirmed, documented rights to use that data for those purposes. Platform terms of service often prohibit scraping. Data licenses often restrict downstream use. What was permissible as raw data may not be permissible as model training material.

The legal landscape here is actively contested. Multiple pending lawsuits in the United States and Europe allege copyright infringement based on the use of authored works in AI training. The outcomes of those cases will affect companies that did not produce the original training data but rely on models that were trained on it—including through third-party APIs where the underlying model's training provenance is unknown.

Platform and API Agreements

Companies accessing AI capabilities through platforms—OpenAI, Anthropic, Google, Microsoft, and their equivalents—are subject to agreements that govern what happens to the data they send. Those agreements have evolved substantially over the past two years, and current terms frequently include provisions that allow the platform to use submitted data for model improvement unless the user has opted out through a mechanism that is not visible in the default product flow.

A company that sends customer data, proprietary information, or regulated data through an AI API without reviewing the current data handling provisions of that API's agreement has made a legal representation it may not be able to keep—specifically, any representation in its own privacy policy or customer contracts about how customer data is used and shared.

User-Generated and Customer Data

Companies that train, fine-tune, or configure AI models using customer data face a distinct layer of risk. The question is whether the company had the right to use that data for that purpose. Most privacy policies and customer agreements in use today were drafted before AI training was a relevant concept. They address data use in terms of service delivery, analytics, and product improvement—categories that courts and regulators are now being asked to interpret in the context of AI training, with results that are not yet settled.

For companies in regulated industries—healthcare, finance, education, legal services—the exposure is heightened. HIPAA, GLBA, FERPA, and their state-level equivalents create specific restrictions on data use that do not contain implied exceptions for AI development.

Common Failure Points in Data Architecture

- Training data obtained through scraping without confirmed license rights
- API or platform data handling terms not reviewed against current version
- Customer data used to fine-tune models without explicit contractual authorization
- No documentation of data provenance for models in production use
- Privacy policy data use representations inconsistent with actual AI data flows
- Regulated data (PHI, PII, financial records) in AI pipelines without sector-specific analysis
- No vendor data processing agreements with AI subprocessors

What Structured Data Architecture Looks Like

A company with sound data architecture for its AI operations can answer the following questions without hesitation: Where did the training data come from, and what license rights does the company hold for each source? What do the platform agreements governing API access currently say about data submitted for inference? Does the company's privacy policy accurately describe how customer data interacts with its AI systems? Are there data processing agreements in place with every vendor that touches input or output data? Has the company conducted a data inventory sufficient to identify regulated data categories in its AI pipelines?

That documentation does not need to be voluminous. It needs to be accurate, current, and defensible.

Data Architecture Self-Assessment

☐	We have documented the sources of all training or fine-tuning data used in our AI systems.
☐	We have confirmed that our license rights for each data source permit its use for AI training or fine-tuning.
☐	We have reviewed the current data handling and training opt-out provisions of every AI platform API we use.
☐	Our privacy policy accurately describes how customer data interacts with our AI systems, including any third-party model APIs.
☐	We have data processing agreements in place with every vendor that handles our input or output data.
☐	We have identified any regulated data categories (PHI, PII, financial records, student records) in our AI data pipelines and applied applicable compliance frameworks.
☐	We have procedures to respond to data subject access requests that intersect with AI-processed data.
☐	We have documentation sufficient to answer data provenance questions during investor due diligence.

LAYER R Risk Allocation

When an AI Output Causes Harm, Who Pays?

Risk allocation is the question every commercial relationship involving AI must answer, and almost none of them do. The contracts that govern AI-related transactions—platform agreements, vendor contracts, commercial agreements, employment arrangements—were largely drafted before the specific risk profile of AI was understood. They address liability in terms of the relationships those contracts were designed for: software delivery, professional services, data licensing. They do not address the novel risks that AI introduces.

The result is a liability landscape where significant risk sits unallocated. When an AI output causes a customer to make a bad decision, generates a legal claim, exposes sensitive data, or produces content that creates reputational or regulatory exposure—the question of who bears the cost is not answered by the contracts in place. It is answered by whoever has the resources and the incentive to litigate it.

The Three Risk Allocation Gaps

Gap 1: Platform Agreements Disclaim Everything

The standard terms of service for major AI platforms—covering API access, model outputs, and data processing—disclaim liability for AI-generated outputs with provisions broad enough to cover nearly any harm the output might cause. The platforms are not wrong to do this. Their outputs are probabilistic, context-dependent, and used in downstream applications they do not control. But the downstream company—the one deploying those outputs in a product or workflow—has not disclaimed that liability to its own customers. It has absorbed it.

A company that deploys a model's outputs without negotiating indemnification rights from the platform, and without allocating output liability in its own customer agreements, is standing between two sets of unremediated risk. It cannot recover upstream from the platform. It has not allocated downstream to the customer. It holds everything in the middle.

Gap 2: Commercial Agreements Do Not Reflect AI Involvement

When a company uses AI to assist in delivering a service, creating content, generating analysis, or making decisions—and does not disclose or contractually address that AI involvement—it has created a representation gap. Customer contracts typically contain representations about how services are performed, what standards apply, and what the customer can rely on. Introducing AI into a workflow that the contract describes in human-performance terms creates a potential breach at the moment the first AI output is used.

Beyond disclosure, the allocation question remains: if an AI-assisted output is wrong, who bears the cost? The contract answer, in most cases, is silence. Courts and arbitrators filling that silence will not always fill it in the company's favor.

Gap 3: Indemnification Chains Are Broken

Enterprise commercial agreements frequently contain indemnification chains: vendor indemnifies the company, company indemnifies the customer. Those chains were designed for identifiable, human-made errors with traceable causes. AI errors are often neither. A model output that causes harm may not be traceable to a specific design decision, training choice, or deployment parameter

in a way that maps onto standard indemnification frameworks. The chain may exist in the contract and be unenforceable in practice.

Questions Every AI-Integrated Company Must Be Able to Answer

If our AI system produces an output that causes financial or reputational harm to a customer, what does our commercial agreement say about who bears that cost?

Do our vendor contracts contain any indemnification provisions specific to AI outputs? If so, are those provisions actually enforceable against the vendor's standard disclaimers?

Does our product liability insurance cover AI-generated output claims? Has our insurer been informed that AI outputs are part of our product delivery?

If an AI output is used in a regulated context—a medical recommendation, a financial decision, a legal document—do our agreements contain the human oversight provisions that applicable law requires?

Have we disclosed AI use in our commercial agreements at a level sufficient to prevent a later argument that the customer did not understand what they were purchasing?

Structuring AI Risk in Contracts

Addressing risk allocation in AI-integrated businesses requires changes to documents in three places: platform and vendor agreements (negotiating AI-specific indemnification and liability provisions), commercial agreements (defining AI involvement, limiting liability for AI outputs, and establishing appropriate disclaimers), and insurance (confirming coverage with carriers and adjusting policies to reflect AI-generated output risk).

None of these changes requires prohibitively complex documentation. What they require is intention. The default terms, in every category, were not designed for this risk profile. They must be deliberately revised.

Risk Allocation Self-Assessment

☐	We have reviewed the indemnification and liability provisions in our primary AI platform and vendor agreements within the past 12 months.
☐	Our commercial agreements with customers include provisions that address AI involvement in service delivery, output accuracy, and liability for AI-generated content.
☐	We have evaluated whether our commercial agreements' representations about service delivery remain accurate in light of AI integration.
☐	We have confirmed with our insurance carrier that our current policies cover claims arising from AI-generated outputs.
☐	We have mapped the indemnification chain across our AI vendor stack and identified any gaps between our upstream protections and downstream obligations.

☐	We have limitation of liability provisions in our commercial agreements that are calibrated to the risk profile of AI-assisted delivery.
☐	We have a defined escalation path for AI-output disputes that does not default to uncapped liability.

LAYER I IP and Output Ownership

Who Owns What the AI Created—and Can That Ownership Be Enforced?

Intellectual property ownership in AI-generated work is one of the least-settled areas of law in the current commercial landscape. The U.S. Copyright Office has taken the position that works generated by AI without sufficient human authorship are not eligible for copyright protection. Courts have affirmed this position in initial cases. The implication—that AI-generated outputs may not be protectable as intellectual property—has not yet been fully absorbed by the companies building products around those outputs.

The problem compounds across the value chain. A company creates AI-generated content. It delivers that content to a customer under a contract that includes an IP assignment. The customer believes it has received protectable intellectual property. The company has represented, at least implicitly, that it can transfer what it is conveying. If the content is not copyrightable—because the human contribution was insufficient—neither party has what they thought they had.

The IP Questions No Standard Agreement Answers

Ownership of AI Outputs

Current U.S. copyright doctrine requires human authorship for copyright protection. Works generated entirely by AI are not eligible. Works involving meaningful human creative contribution—selecting, arranging, modifying AI outputs in ways that reflect human judgment—may be eligible, but the standard is not precise, and the Copyright Office reviews AI-involved works on a case-by-case basis.

The practical consequence: companies that deliver AI-generated content as a product should not represent that the content is automatically eligible for copyright protection. They should document the human creative contribution that exists, maintain records of that contribution, and structure their commercial agreements to reflect the actual ownership position—not an assumed one.

Platform License-Back Provisions

Most AI platform agreements include provisions that grant the platform a license to outputs generated using their models. The scope of those provisions varies significantly. Some are narrow (limited to improving the model). Some are broad (including commercial use rights to derivative works). Some change with platform updates without requiring affirmative customer acceptance.

A company that has granted a broad license-back to its AI platform may not be in a position to grant an exclusive IP license to a customer—even if the customer's agreement purports to convey one. The conflict between the platform's retained rights and the customer's purported ownership is not hypothetical. It is a structural defect in the IP chain that exists in any company that has not audited its platform agreements against its commercial IP representations.

Contractor and Employee AI Work Product

Standard IP assignment agreements in employment contracts and independent contractor arrangements were drafted to address human-created work product. They typically assign to the company all work product created within the scope of employment or the services arrangement. Whether AI-generated outputs created by an employee or contractor in the course of their work fall within those assignment provisions is not settled. The answer depends on the specific language of the agreement, the jurisdiction, and the nature of the AI contribution.

Companies that rely on AI-generated work product created by employees or contractors should review their existing IP assignment provisions to confirm they are adequate—and update them where they are not.

Open-Source Model Compliance

Companies that use open-source AI models—whether as base models for fine-tuning, as inference engines, or as components of a larger system—are subject to the license terms governing those models. Open-source AI licenses vary significantly. Some are permissive (MIT, Apache 2.0). Some contain copyleft provisions that require derivative works to be distributed under the same license. Some AI-specific licenses (including the RAIL license family) contain use restrictions that prohibit specific commercial applications regardless of whether the licensee modified the model.

A company that has built a commercial product on an open-source model without auditing the applicable license is operating with an unknown constraint on its IP position, its ability to raise capital, and its ability to be acquired.

The IP Audit Every AI Company Needs

- Identify every AI model component in your product or workflow and its applicable license
- Review platform agreements for license-back provisions and scope of retained rights
- Assess open-source license obligations and confirm compliance
- Evaluate copyright eligibility of AI-generated outputs delivered to customers
- Review employment and contractor IP assignments for AI work product coverage
- Update commercial agreements to accurately reflect IP ownership positions
- Document human creative contributions to AI-generated outputs where copyright is asserted

IP and Output Ownership Self-Assessment

☐	We have identified the license terms governing every AI model component used in our product or operations.
☐	We have reviewed our primary platform agreements for license-back provisions and confirmed they do not conflict with IP representations in our customer agreements.
☐	We have evaluated the copyright eligibility of AI-generated outputs that we deliver to customers or claim as proprietary.
☐	Our commercial agreements accurately reflect our IP ownership position with respect to AI-generated content—including any limitations on copyright protection.
☐	Our employment and contractor IP assignment provisions have been reviewed and updated to address AI-generated work product.
☐	We have audited our use of open-source AI models for license compliance, including copyleft and use restriction provisions.
☐	We have a documented process for maintaining records of human creative contribution to AI-generated outputs where copyright protection is asserted.

LAYER A_2 Accountability and Governance

Can You Prove You Were Careful?

In the event of an AI-related incident—a harmful output, a regulatory inquiry, a litigation claim—the first question any competent adversary will ask is not what the AI did. It is what the company did to prevent, detect, and respond to harm. Documentation of governance and oversight is not just a regulatory compliance item. It is a litigation posture. It is a due diligence data room requirement. It is the difference between demonstrating organizational competence and demonstrating organizational negligence.

Most early- and growth-stage companies have no AI governance documentation. Not because leadership has decided the risk is acceptable, but because governance was not part of the original product roadmap, and no one has yet assigned it a place on the operational agenda.

That gap closes voluntarily—through structured governance documentation—or involuntarily, through the response to an event that required governance that did not exist.

What AI Governance Requires

AI Use Policy

An AI use policy is a formal internal document that defines how employees and contractors may use AI tools in their work. It addresses which tools are approved for which purposes, what categories of data may be submitted to AI systems, what review processes apply to AI-generated outputs before they are used in client-facing or consequential contexts, and what reporting requirements exist when an AI tool produces unexpected or potentially harmful results.

The absence of an AI use policy does not mean employees are not using AI. It means they are using it without guidance, without oversight, and without documentation that the company took any responsibility for managing the risk. In a dispute over an AI-related harm, that absence will be noted.

Human Oversight Protocols

Governance frameworks increasingly require—and courts and regulators are beginning to expect—evidence that AI outputs in consequential contexts were subject to meaningful human review before being acted upon. This is particularly true in regulated industries (healthcare, finance, law, education) but is becoming relevant across business contexts as AI outputs are used in hiring, underwriting, credit decisions, content moderation, and customer-facing advice.

Human oversight documentation does not require that every AI output be reviewed by a human. It requires that the company has defined where human review is required, what that review consists of, and who is responsible for it. The policy needs to be specific enough to be followed and documented enough to be proven.

Incident Response

AI incident response is a defined, documented protocol for identifying, containing, and responding to AI-related failures. An incident can take many forms: a model output that is factually wrong in a material context, a data incident arising from AI data processing, a third-party claim arising from AI-generated content, or a regulatory inquiry related to AI use.

A company without an incident response protocol is not a company that will never have an incident. It is a company that will improvise when one occurs. Improvised responses to AI incidents produce inconsistent fact patterns, poor documentation, and—in regulatory and litigation contexts—the appearance of disorganization that is treated as evidence of systemic failure rather than an isolated event.

Board-Level AI Risk Oversight

Institutional investors, insurance underwriters, and enterprise customers are increasingly asking for evidence that AI risk is managed at the board or senior leadership level—not just operationally. A board-level AI risk summary, updated at least annually, demonstrates that the company takes AI risk seriously enough to govern it at the appropriate level of organizational authority.

This document does not need to be lengthy. It needs to accurately summarize the company's AI risk profile, the governance measures in place, and the open questions the company is actively managing. Its existence demonstrates organizational maturity. Its absence, in a diligence context, raises questions that are difficult to answer favorably after the fact.

The Governance Documentation Stack

Every AI-integrated company should maintain, at minimum, the following documents:

- AI Use Policy — approved tools, data handling requirements, output review standards
- Human Oversight Protocol — where review is required, by whom, and how it is documented
- AI Incident Response Procedure — identification, containment, notification, and remediation
- Vendor AI Risk Assessment — evaluation of third-party AI tools against defined risk criteria
- Board-Level AI Risk Summary — annual summary of AI risk profile and governance status
- Employee Training Record — documentation of AI policy training and acknowledgment

Accountability and Governance Self-Assessment

☐	We have a written AI Use Policy that has been communicated to all employees and contractors.
☐	Our AI Use Policy identifies which tools are approved, what data restrictions apply, and what review standards govern AI-generated outputs.
☐	We have defined, in writing, where human review of AI outputs is required before those outputs are used in consequential contexts.
☐	We have an AI Incident Response Procedure that identifies how AI-related failures are reported, documented, and escalated.
☐	We have conducted a documented risk assessment of the third-party AI tools and platforms integrated into our operations.
☐	We maintain records of employee acknowledgment of AI-related policies.

☐	We have presented a summary of our AI risk profile and governance approach to the board or senior leadership within the past 12 months.
☐	Our governance documentation is current and would withstand review in a due diligence process or regulatory inquiry.

LAYER 2 Regulatory Architecture

Which Legal Frameworks Govern Your AI Operations—Today and in the Coming 24 Months?

AI regulation is forming in multiple jurisdictions simultaneously, at different speeds, with different scopes, and with varying degrees of enforceability. The companies that will be best positioned when regulatory frameworks consolidate are those that have already assessed their exposure—not those that begin their compliance work when the first enforcement action occurs.

The current regulatory landscape for AI is not empty. Multiple frameworks are in effect now. Others are in the process of implementation. The absence of a single, unified U.S. federal AI regulation does not mean regulatory risk is absent. It means that risk is fragmented across sectoral regulators, state laws, and existing legal doctrines—which makes it harder to assess and easier to miss.

Frameworks in Effect Now

U.S. Federal Sectoral Regulation

Federal agencies with existing authority over specific industries have extended that authority to AI-related conduct without waiting for comprehensive AI legislation. The FTC has issued guidance and brought enforcement actions under Section 5 of the FTC Act addressing deceptive AI claims and AI-related privacy violations. The EEOC has issued guidance on employer liability for discriminatory AI use in hiring and employment decisions. The CFPB has taken the position that existing fair lending and fair credit laws apply to AI-driven credit decisions. The FDA has issued guidance on AI/ML-based software as a medical device. These are not future regulatory developments. They are current enforcement priorities.

U.S. State AI and Privacy Laws

Approximately a dozen U.S. states have enacted laws with specific AI provisions, and several more are in active legislative sessions. Colorado's AI Act imposes obligations on developers and deployers of high-risk AI systems. Illinois's Artificial Intelligence Video Interview Act regulates AI use in employment contexts. Several states have enacted or proposed laws governing AI-generated deepfakes, synthetic media, and automated decision-making in consequential contexts.

For companies operating nationally, the patchwork of state laws creates a compliance obligation that cannot be addressed by reference to a single standard. The appropriate response is a jurisdiction-mapping exercise that identifies which state laws apply to which aspects of operations—not a general presumption that federal standards are sufficient.

GDPR and EU AI Act

U.S.-based companies with EU customers, employees, or data subjects are subject to GDPR's restrictions on automated decision-making and profiling. Article 22 of GDPR restricts decisions based solely on automated processing that have legal or similarly significant effects on individuals, and creates a right to human review of such decisions. Many AI applications in use today are within the scope of this provision, and many companies operating under it have not assessed compliance.

The EU AI Act entered into force in August 2024, with phased implementation running through 2027. It applies to AI systems placed on the EU market or used in the EU, regardless of where the developer is located. It establishes a risk-tiered framework with prohibitions on certain high-risk applications, mandatory obligations for providers and deployers of high-risk systems, and

transparency requirements across a broader category of AI applications. U.S. companies with EU exposure should be evaluating their AI applications against this framework now, not when implementation deadlines arrive.

Sector-Specific Frameworks

Healthcare companies using AI are subject to HIPAA's restrictions on AI data processing and the FDA's guidance on AI/ML-based software in clinical contexts. Financial services companies using AI in credit, insurance, or investment decisions face obligations under the Equal Credit Opportunity Act, the Fair Housing Act, and relevant prudential regulatory guidance. Companies using AI in education face FERPA and COPPA considerations. Companies using AI in employment face EEOC guidance and state-level employment non-discrimination laws that apply to algorithmic selection tools.

These are not theoretical future obligations. They are frameworks that regulators are actively applying to AI conduct today.

Regulatory Horizon — Key Dates and Developments

EU AI Act — Prohibition provisions effective February 2025. High-risk system obligations begin August 2026.

U.S. State Laws — Multiple state AI disclosure, employment, and automated decision laws effective in 2024-2026. Jurisdiction mapping is required for any nationally operating company.

FTC AI Guidance — Ongoing enforcement posture under existing authority. Companies making AI-related claims in marketing are subject to active FTC scrutiny.

CFPB Automated Decision Guidance — Financial services companies using AI in credit, collections, or customer service face expanding examination scrutiny.

Federal Legislation — Multiple AI bills in various stages of the legislative process. No single comprehensive federal law has passed as of this writing, but the regulatory direction is established.

How to Assess Your Regulatory Exposure

Regulatory assessment for AI does not require predicting legislative outcomes. It requires mapping your current AI applications against the frameworks already in effect and identifying where your operations intersect with regulated domains. That mapping exercise—done properly—identifies the highest-priority compliance actions and builds the evidentiary foundation for demonstrating good-faith compliance if a regulatory inquiry occurs.

The alternative—assuming that because no enforcement action has landed yet, none will—is not a risk posture. It is a deferred compliance obligation that accumulates interest.

Regulatory Architecture Self-Assessment

☐	We have identified the federal sectoral regulatory frameworks that apply to our AI use cases (FTC, EEOC, CFPB, FDA, or others applicable to our industry).
☐	We have mapped which U.S. state AI and privacy laws apply to our operations and confirmed compliance with the provisions now in effect.
☐	We have assessed whether GDPR's Article 22 automated decision-making restrictions apply to any of our AI applications.
☐	We have evaluated our AI product or service against the EU AI Act's risk classification framework and identified any high-risk system obligations.
☐	We have assigned regulatory monitoring responsibility to a specific person or function within the organization.
☐	We have documented our regulatory compliance posture for AI in a format that can be reviewed by counsel, auditors, or regulators.
☐	We have a process for updating our compliance assessment when new regulations are enacted or when our AI applications change materially.

Self-Assessment Summary and Next Steps

You have now worked through all five layers of the ARIA Framework. Before identifying next steps, consolidate your findings. For each layer, assess your current status across four categories:

ARIA Layer	Current Status	Priority	Immediate Action
A — Data Architecture	*[ ] Structured [ ] Partial [ ] None*	*Critical / High / Medium*	_____
R — Risk Allocation	*[ ] Structured [ ] Partial [ ] None*	*Critical / High / Medium*	_____
I — IP & Output Ownership	*[ ] Structured [ ] Partial [ ] None*	*Critical / High / Medium*	_____
A2 — Governance	*[ ] Structured [ ] Partial [ ] None*	*Critical / High / Medium*	_____
2 — Regulatory Architecture	*[ ] Structured [ ] Partial [ ] None*	*Critical / High / Medium*	_____

Interpreting Your Results

If you have identified one or more layers where your current status is None or Partial, and where the priority level is Critical or High, your company has addressable legal exposure that warrants structured attention before your next funding round, transaction, or significant regulatory development.

That statement is not a prediction of harm. It is an observation about risk structure. Companies with unaddressed AI legal exposure do not invariably suffer consequences from it. But companies with addressed exposure are in a materially better position when investors conduct diligence, when customers request contractual representations, when regulators inquire, or when a dispute arises that requires the company to demonstrate it acted reasonably.

The cost of addressing that exposure in advance is fixed and finite. The cost of addressing it reactively is not.

What to Do Next

If you have gaps in one or two layers:

Begin with a targeted review of the specific documents and agreements implicated by those gaps. For Layer A issues: review your current platform agreements and data processing documentation. For Layer R issues: review your primary vendor agreements and commercial contract templates. For Layer I issues: audit your AI model licenses and commercial IP representations. For Layer A2 issues: draft an AI Use Policy and Human Oversight Protocol. For Layer 2 issues: commission a regulatory mapping exercise.

If you have gaps across three or more layers:

The ARIA Certification is the appropriate next step. Addressing individual layers in isolation can close one exposure while missing the systemic gaps that only a structured cross-layer assessment

reveals. The ARIA Certification provides a formal, scored credential you can present to investors, customers, and your board.

The ARIA Certification

AI Legal Readiness Score | Issued by Oak & Hill, P.C.

The AI legal risk categories covered in this guide are real, measurable, and increasingly consequential. Investors are asking about them. Enterprise customers are requiring contractual representations about them. Regulators are enforcing against companies that ignored them.

Until now, there has been no standardized way for a company to demonstrate formally and credibly that its AI legal posture has been professionally evaluated. The ARIA Certification was built to fill that gap.

What the ARIA Certification Is

The ARIA Certification is a formal, scored AI legal readiness assessment issued by Oak & Hill, P.C. as a professional credential. It is attorney-issued, documented for third-party reliance, and renewed annually.

Think of it as what SOC 2 did for data security, or what a 409A valuation does for equity pricing. SOC 2 converted an abstract question about data security into a verified, standardized answer that investors, customers, and auditors could rely on. The 409A converted a contested equity valuation question into a defensible, recognized standard the IRS, boards, and employees accepted. The ARIA Certification does the same thing for AI legal readiness: it converts the question every investor and acquirer is now asking into a formal, scored, attorney-issued answer they can point to.

Standard	What It Answers	Who Relies on It
SOC 2	*Is this company's data handling secure and audited?*	Enterprise customers, investors, insurers
409A Valuation	*What is the fair market value of common stock?*	IRS, boards, employees, investors
ARIA Certification	*Is this company's AI legal infrastructure formally structured?*	Investors, acquirers, customers, insurers, regulators

The Five Deliverables

Every ARIA Certification is conducted by a licensed attorney at Oak & Hill, P.C. and produces five formal deliverables:

#	Deliverable	Description
1	**ARIA Score Summary**	One-page scored dashboard across all five layers. Formatted for investor data rooms and due diligence packages.
2	**ARIA Certification Letter**	Formal letter on Oak & Hill letterhead confirming assessment scope, composite score, issuance date, and issuing attorney. This is the reliance document.

3	**ARIA Detailed Report**	10-15 page analysis organized by layer, documenting findings, exposure levels, and the basis for each score. Attorney work product.
4	**ARIA Remediation Roadmap**	Prioritized, sequenced action list with estimated effort and urgency level for each gap identified.
5	**ARIA Digital Seal**	Formal certification seal for investor materials, RFP responses, and marketing. Includes issue date and 12-month validity.

The ARIA Score

Each layer is scored 1-4 by the assessing attorney. The composite average produces a letter grade and determines whether a full Certification or an Assessment Certificate is issued.

Grade	Score	Credential	Third-Party Reliance Status
A	4.0-5.0	**ARIA Certified — Distinguished**	Suitable for institutional diligence, M&A, and enterprise procurement without qualification.
B	3.0-3.9	**ARIA Certified — Proficient**	Appropriate for most investor and customer contexts with disclosed minor gaps.
C	2.0-2.9	**ARIA Assessed — Developing**	Material gaps documented. Not reliance-grade without disclosed remediation plan.
D	Below 2.0	**ARIA Assessed — Early Stage**	Certification withheld. Full assessment and remediation roadmap issued.

Who Should Obtain an ARIA Certification

- Any company approaching a Series A or B raise where AI is material to the product or operations
- Any company in an M&A process, whether as seller or as target
- Any company entering an enterprise sales cycle with customers requiring AI compliance representations
- Any company applying for D&O or cyber insurance that asks about AI risk governance
- Any company whose investors or board have begun asking about AI legal posture
- Any company that has completed this Playbook and identified gaps in two or more ARIA layers

Get Your ARIA Certification

The first standardized, attorney-issued AI legal readiness credential.

$5,000 flat fee | 3-4 week timeline | 12-month validity

Oak & Hill, P.C.
oakhillcounsel.com
Attorneys at Law | AI Legal Infrastructure Practice