RIZE™ TECHNOLOGIES

AI Readiness for Law Firms:
A Five-Part Series for IT Leaders

PART 2 – SECURITY AND CONFIDENTIALITY CONTROLS

2

www.rizetechnologies.com

Introduction to Ebook Series

The legal industry stands at an inflection point. Artificial intelligence has moved from theoretical promise to practical reality, with tools that can draft contracts, conduct legal research, analyze discovery documents, and predict case outcomes. For mid-size law firms, the question is no longer whether to adopt AI, but how to do so responsibly, securely, and in compliance with ethical obligations.

While firm leadership focuses on strategic AI adoption and attorneys explore use cases, IT teams face a critical challenge: building the technical foundation that makes safe, effective AI deployment possible. Without proper preparation, firms risk data breaches, ethical violations, regulatory non-compliance, and malpractice exposure.

This five-part series provides IT leaders at mid-size law firms with a practical roadmap for AI readiness, covering the technical, security, and operational considerations essential for successful AI integration. If law firms lack the appropriate IT resources on their staffs, the professional services team at Rize Technologies is available to help them plan and implement a customized roadmap for AI readiness.

Security and Confidentiality Controls

Security failures in AI deployment can have catastrophic consequences for law firms. A robust security framework must address both traditional cybersecurity concerns and AI-specific vulnerabilities that many IT teams haven't encountered before.

Security failures in AI deployment can have catastrophic consequences for law firms. Unlike traditional software that processes data in predictable, controlled ways, AI systems involve complex interactions between your firm's data, vendor platforms, and sophisticated models that may be shared across multiple customers.

A robust security framework must address both traditional cybersecurity concerns—such as access controls, encryption, and network security—and AI-specific vulnerabilities that many IT teams haven't encountered before. These include risks like prompt injection attacks where malicious inputs manipulate AI behavior, data leakage through model training where your confidential information could inadvertently become part of a vendor's broader AI capabilities, and the challenge of auditing AI decision-making processes that lack the transparency of conventional software.

The stakes are particularly high because AI tools often require **access to large volumes of sensitive data to deliver value,** creating concentrated risk that demands equally concentrated protective measures.

Implement Rigorous Access Controls

Directly related to the security issues summarized above, AI tools must respect the same confidentiality boundaries that govern human access to client information. This means deploying role-based access controls that align AI tool permissions with matter-level access rights. An associate working on a securities litigation shouldn't be able to use an AI tool to access documents from an unrelated family law matter, even if both matters exist in the same document management system.

The best technical implementation requires multi-factor authentication for all AI platform access, isolated environments for different practice groups or sensitive matters, and comprehensive audit trails. Those audit trails need to capture more than traditional systems do—not just who accessed what data, but what prompts they used, what outputs the AI generated, and when those outputs were shared or incorporated into work product. You also need session timeouts and automatic logoffs to prevent unauthorized access through unattended workstations.

Some firms find it valuable to implement an AI access request workflow where attorneys must justify business need before gaining access to specific AI capabilities, particularly for tools that process sensitive client data. This might seem bureaucratic, but it creates a moment of reflection that helps prevent inappropriate AI use before it happens.

Conduct Thorough Vendor Due Diligence

When evaluating AI vendors, your standard technology vetting process is likely insufficient. You need to dig deeper into questions that might not matter for traditional software but are critical for AI. Where exactly is data stored, and is it encrypted both at rest and in transit? How long is data retained after processing? Is firm data used to train models that other customers can access? These questions often reveal deal-breaking issues—a vendor whose terms allow them to train their public model on your client data is fundamentally incompatible with your confidentiality obligations.

Security certifications matter but dig into what they really mean. A SOC 2 Type II certification is valuable, but you need to understand the scope and any exceptions. Ask about sub-processors and dependencies—what third parties have access to firm data, and are they subject to equivalent security standards? The AI vendor might have excellent security, but if they rely on a cloud storage provider with weaker controls, your data is only as secure as the weakest link.

Incident response capabilities deserve special attention. What is the vendor's breach notification timeline? What support do they provide during security incidents? Can they help you meet your obligations to notify clients, bar authorities, and regulators if something goes wrong? These questions reveal whether a vendor treats security as a compliance checkbox or a genuine commitment.

Finally, review contracts carefully for provisions around data ownership, portability rights, deletion procedures, and liability limitations. The standard terms you'll receive often favor the vendor heavily. Negotiate stronger protections where possible, particularly regarding use of firm data for model training. Some vendors will agree to contractual terms prohibiting training on customer data if you ask, but they likely won't offer those terms without negotiation.

Deploy Data Loss Prevention Measures

Even with proper policies, human error remains a significant risk. Attorneys working under deadline pressure might not stop to consider whether the document they're uploading to an AI tool contains information that shouldn't leave the firm's secure environment. Technical controls provide a safety net.

Configure data loss prevention tools to monitor and restrict sensitive information being submitted to AI platforms. Implement content filtering that detects and blocks client names, case numbers, social security numbers, and other identifiers. Create protocols requiring document sanitation before AI processing—removing metadata, redacting identifiable information, and stripping formatting that might leak information about firm systems or practices.

For matters involving particularly sensitive information—national security cases, high-profile corporate transactions, matters under protective orders—consider on-premises or private cloud AI deployments that never expose data to external networks. Yes, these solutions cost more and offer fewer features than cutting-edge cloud AI, but stronger security is worthwhile for your most sensitive work.

Establish alerts when large volumes of data are uploaded to AI platforms or when unusual usage patterns emerge. An attorney who suddenly uploads fifty gigabytes to an AI service deserves a friendly check-in call to make sure they understand the implications and are following proper procedures.

The key is balancing security with usability. Overly restrictive controls that impede legitimate work will **drive attorneys toward unauthorized "shadow AI" solutions that circumvent security entirely.** Your goal is making approved tools sufficiently flexible that unauthorized alternatives become unnecessary.

In Part 3 of this 5-part series, we will look at specific network and computing infrastructure issues and requirements as they relate to AI-readiness. Be sure to download the entire series by visiting the Resources section on **rizetechnologies.com**

About This Guide

This guide was developed based on emerging best practices in legal technology, cybersecurity, and professional responsibility for AI adoption in law firms, but it does not constitute legal, technical, or professional advice. While we always strive to provide accurate information, we make no warranties regarding completeness or accuracy, and we assume no liability for any damages, losses, or consequences arising from your use of or reliance on this information. Each firm is solely responsible for evaluating the applicability of this guidance to their situation and for all decisions and actions taken.

About Rize Technologies

Rize Technologies helps law firms thrive in today's fast-moving, AI-powered world. With more than a decade of experience serving the legal industry, we deliver end-to-end IT solutions that keep firms secure, productive, and ready for what's next. From 24/7 managed IT and cloud services to advanced cybersecurity, data protection, disaster recovery, and AI readiness, every solution is built around the unique needs of each firm. Our white-glove support and proactive approach to IT helps minimize downtime and maximize confidence. Discover how your firm can rise higher at

rizetechnologies.com