

FOREIGN TRAINED LAWYERS ASSOCIATION

Policy Paper No. III · April 2026

FTLA POLICY PAPER SERIES
NEW YORK · WASHINGTON · GLOBAL
LEGAL WORKFORCE · ARTIFICIAL INTELLIGENCE ·
PROFESSIONAL INTEGRATION

Artificial Intelligence as Legal Workforce Infrastructure

*Integrating Foreign-Trained Lawyers into the U.S. Economy through AI-Driven Legal Readiness
Systems*

AN FTLA POLICY FRAMEWORK

Sopiko Gugava

Foreign Trained Lawyers Association

*An institutional analysis of artificial intelligence as enabling
infrastructure for cross-border legal workforce integration.*

FOREIGN TRAINED LAWYERS ASSOCIATION
NEW YORK, NEW YORK

© 2026 Foreign Trained Lawyers Association

**M M X X V I
III**

ABSTRACT

The integration of foreign-trained lawyers into the U.S. legal profession is constrained, as FTLA Policy Paper No. 1 documented, by structural friction in admission pathways, by the absence of organized professional development infrastructure, and by a pass rate gap on the U.S. bar examination that reflects, in significant part, the structural difficulty of adapting from civil-law or other foreign legal training to the format and substantive expectations of a U.S. multistate examination. This paper takes the position that artificial intelligence, properly deployed as workforce-integration infrastructure rather than as a substitute for legal labor, offers a measured response to each of these constraints, and that the United States has a national-interest reason to develop and deploy AI-driven legal readiness systems that compress the integration timeline for foreign-trained lawyers without lowering substantive professional standards.

The paper develops this position in three steps. It identifies the specific integration tasks that AI is well suited to support, including comparative-law translation, bar examination preparation calibrated to foreign-trained candidate profiles, regulatory-context onboarding for U.S. specialty practice, and ethics infrastructure for cross-border practice. It situates these tasks within the broader U.S. policy framework on AI in legal services, including the American Bar Association's Formal Opinion 512 (July 2024) on the use of generative AI by lawyers, the National Institute of Standards and Technology AI Risk Management Framework, and emerging state bar guidance. And it sets out a four-component policy framework through which FTLA, in coordination with state bar admission authorities, ABA-accredited LL.M. programs, and federal economic-development apparatus, can deploy AI-driven legal readiness infrastructure on terms consistent with professional responsibility, equitable access, and the U.S. national interest in a more efficiently integrated cross-border legal workforce.

TABLE OF CONTENTS

Abstract.....	i
Glossary of Abbreviations	iii
I. Introduction: The Workforce-Integration Frame.....	1
A. The continuing capacity gap	2
B. Why AI is the right tool for this problem.....	4
II. The Problem Statement	6
A. The structural friction in foreign-trained lawyer integration	6
B. The bar examination performance gap	8
C. The post-admission integration gap	10
III. The Existing Framework.....	12
A. ABA Formal Opinion 512 and the duties of competence and confidentiality	12
B. State bar guidance on generative AI.....	15
C. The NIST AI Risk Management Framework.....	17
D. Federal workforce and immigration architecture	19
IV. Analysis	21
A. Four legal-readiness functions AI can perform well.....	21
B. Four functions AI cannot perform and should not be assigned.....	26
C. The equity and access dimension	28
D. Economic stakes: the FDI legal services connection	30
V. Comparative Perspective.....	32
A. The European Union: the AI Act and legal-services applications	32
B. The United Kingdom: SRA guidance and the regulatory sandbox	34
C. Singapore and Canada: targeted use cases	36
VI. Policy Recommendations	38
A. The FTLA Legal Readiness Platform	38
B. AI-assisted bar preparation calibrated to foreign-trained candidates	40
C. Cross-border practice ethics infrastructure.....	42
D. Coordination with federal workforce and FDI agencies	44
VII. Implementation Considerations.....	46
VIII. Conclusion	49
References.....	51

GLOSSARY OF ABBREVIATIONS

The following abbreviations are used throughout this paper. Citations to source materials appear in the References.

Term	Full term
ABA	American Bar Association
AI	Artificial Intelligence
AI Act	Regulation (EU) 2024/1689 on Artificial Intelligence
BOLE	Board of Law Examiners (used herein principally for New York)
CLE	Continuing Legal Education
DOL	U.S. Department of Labor
FDI	Foreign Direct Investment
FTLA	Foreign Trained Lawyers Association
GenAI	Generative Artificial Intelligence
LL.M.	Master of Laws degree
LLM	Large Language Model (when distinguishable from LL.M.)
MPRE	Multistate Professional Responsibility Examination
MPT	Multistate Performance Test
NCBE	National Conference of Bar Examiners
NIST	National Institute of Standards and Technology
NIST AI RMF	NIST Artificial Intelligence Risk Management Framework
OECD	Organisation for Economic Co-operation and Development
RAG	Retrieval-Augmented Generation
SRA	Solicitors Regulation Authority (United Kingdom)
UBE	Uniform Bar Examination
USCIS	U.S. Citizenship and Immigration Services
WIPO	World Intellectual Property Organization

I. Introduction: The Workforce-Integration Frame

The Foreign Trained Lawyers Association was established, as FTLA Policy Paper No. 1 set out, on the premise that foreign-trained lawyers integrated into U.S. legal practice constitute critical professional infrastructure for the U.S. economy. The paper documented the volume and complexity of cross-border compliance demands generated by federal regulatory architecture spanning the Sarbanes-Oxley Act of 2002, the Foreign Investment Risk Review Modernization Act of 2018, the Uyghur Forced Labor Prevention Act of 2021, and related statutes; it documented the size of inbound foreign direct investment, \$5.71 trillion at year-end 2024, and the population of foreign issuers in U.S. capital markets, more than 1,500 exchange-listed companies; and it documented the structural friction in U.S. licensing pathways that constrains the supply of foreign-trained legal capacity available to meet that demand. The diagnostic was that the United States has, in foreign-trained lawyers, a strategic asset whose institutional infrastructure has not been adequately built.

This paper, the third in the FTLA series, addresses one component of that institutional infrastructure: the technological systems through which foreign-trained lawyers can compress the integration timeline from foreign legal training to effective U.S. practice. The premise is that artificial intelligence, deployed not as a substitute for human legal labor but as workforce-integration infrastructure, can address several of the specific pressure points identified in Paper No. 1, comparative-law translation, calibrated bar examination preparation, regulatory-context onboarding for U.S. specialty practice, and ethics infrastructure for cross-border practice, in ways that traditional CLE programming and informal mentorship cannot match in scale or precision.

The paper does not argue that AI should replace existing professional development infrastructure. It argues that AI should be added to that infrastructure on terms consistent with the professional responsibility framework articulated in the American Bar Association's Formal Opinion 512 (July 2024), the AI risk management discipline articulated in the National Institute of Standards and Technology AI Risk Management Framework (January 2023, with the Generative AI Profile released July 2024), and the equitable-access principles that have characterized U.S. legal-services policy across administrations. The argument is, in the precise sense, an argument about institutional design rather than about technology itself.

A. The continuing capacity gap

The capacity gap documented in FTLA Policy Paper No. 1 has not been resolved in the period since that paper's publication. The cross-border compliance architecture, encompassing capital markets disclosure, national security review, supply chain integrity, sanctions enforcement, and responsible business conduct, continues to generate volumes of regulated activity that exceed what the domestically trained U.S. bar can efficiently handle alone. The supply of foreign-trained lawyers admitted to U.S. practice continues to be constrained by bar examination outcomes, by post-admission integration friction, and by the absence of standardized professional development infrastructure for cross-border practice.

In February 2025, foreign-educated candidates were a majority of all New York bar examination takers (2,225 of 4,090, or 54 percent), and their pass rate was 30 percent against an out-of-state ABA graduate

first-time pass rate of 73 percent. The pass rate gap reflects, on Paper No. 1's analysis, the structural mismatch between standard bar preparation programming, designed for ABA graduates emerging directly from a J.D. program, and the actual learning needs of foreign-trained candidates. The gap is not, in any meaningful sense, a gap in candidate ability; it is a gap in the calibration of preparation infrastructure to candidate profile. AI-driven legal readiness systems are, in the analysis offered here, the most direct mechanism through which that calibration can be achieved at scale.

Beyond bar admission, the post-admission integration friction described in Paper No. 1 persists. Newly admitted foreign-trained lawyers entering U.S. practice in solo, small-firm, or in-house settings face a learning curve on U.S. specialty practice, on local procedural conventions, on client-development practices calibrated to U.S. market norms, and on the substantive law of the regulatory areas in which cross-border practice is concentrated. Traditional CLE programming addresses some of this learning need, but in modular form that does not deliver the integrated, contextualized, asynchronous learning that foreign-trained lawyers entering practice can most effectively use. The result is an integration timeline that, anecdotally, runs three to five years from bar admission to genuine professional fluency, against an integration timeline for domestically trained lawyers that runs approximately one to two years.

B. Why AI is the right tool for this problem

Several features of the workforce-integration problem make it particularly amenable to AI-driven intervention. First, the problem is fundamentally one of personalized learning at scale: each foreign-trained lawyer arrives in U.S. practice with a particular foreign legal training, a particular subset of U.S. law in which their work concentrates, a particular language profile, and a particular set of comparative gaps that need to be closed. Traditional CLE and bar preparation programming cannot cost-effectively deliver personalized learning at the level of granularity that this problem requires; AI-driven systems can.

Second, the problem is fundamentally one of comparative analysis: it requires the learner to identify, for each U.S. doctrinal area, the points of convergence and divergence with the foreign legal system in which the learner was originally trained. AI systems trained on legal corpora can perform comparative analysis, presented as a learning tool to the foreign-trained lawyer, in ways that scale across the diverse source jurisdictions, more than 50 in any given LL.M. cohort, that the U.S. foreign-trained lawyer population represents. No single human instructor, however expert, can reliably perform this comparative analysis across all relevant source jurisdictions at the depth that effective integration requires.

Third, the problem is one for which substantial training data exists. The U.S. legal corpus, statutes, regulations, judicial decisions, ABA Model Rules, NCBE bar examination materials, leading treatises, is large, well-structured, and substantially digitized. The corresponding foreign legal corpora, particularly for the major source jurisdictions of foreign-trained lawyers in the United States, are also substantially digitized and accessible. AI systems can, in principle, be trained on these corpora to produce learning tools calibrated to the comparative-method analytical task that foreign-trained legal integration requires.

Fourth, the problem is one in which the AI system's role is properly understood as supplementary rather than substitutive. The AI system does not practice law, does not give legal advice to clients, and does not

replace the supervising lawyer's professional judgment. It supports the foreign-trained lawyer's own learning and integration trajectory. This supplementary role is the role that ABA Formal Opinion 512 and parallel state guidance have identified as appropriate for generative AI in legal services, and it situates the proposed framework within the existing professional responsibility architecture rather than asking that architecture to expand to accommodate a new use case.

II. The Problem Statement

The foreign-trained lawyer integration problem, framed in workforce-infrastructure terms, comprises three distinct sub-problems. Each is empirically documented in the data presented in FTLA Policy Paper No. 1 and the data updated in this section. Each is, on the analysis offered here, addressable through AI-driven legal readiness systems deployed on terms consistent with the existing professional responsibility framework. This section sets out the three sub-problems in turn.

A. The structural friction in foreign-trained lawyer integration

The structural friction in foreign-trained lawyer integration begins before the bar examination and continues through the early years of U.S. practice. It encompasses the LL.M. application and admission process, the foreign credentials evaluation conducted by the New York State Board of Law Examiners under Rule 520.6 and analogous rules in other jurisdictions, the substantive transition from foreign legal training to U.S. legal training during the LL.M. year, the bar examination preparation process, and the post-admission entry into U.S. practice. Each step in the sequence imposes costs, time, and uncertainty that, in the aggregate, deter qualified foreign-trained candidates from pursuing U.S. integration and lengthen the timeline for those who do.

The friction is most acute at three points. First, the LL.M. year is, for many foreign-trained candidates, the first sustained encounter with the U.S. legal system, and the standard LL.M. curriculum was not designed primarily as bar preparation. Foreign-trained candidates emerging from LL.M. programs frequently report that the substantive bar tested subjects, particularly Constitutional Law, Civil Procedure, Evidence, and the common law foundations of Contracts, Torts, and Property, received limited treatment in their LL.M. coursework. This curricular gap is partly a feature of LL.M. program design, which has historically prioritized comparative and specialized coursework over foundational common-law instruction, and partly a feature of the time available within a one-year program. Whatever its source, the gap translates into bar examination underperformance.

Second, the bar preparation period itself, typically the months immediately following LL.M. completion, is poorly calibrated to foreign-trained candidate needs. The leading commercial bar preparation providers offer comprehensive programs designed for ABA graduates emerging from a three-year J.D. program; the programs assume substantial prior exposure to the bar tested subjects and use the preparation period for review and refinement rather than for foundational instruction. Foreign-trained candidates using these programs frequently report that the program pace is too fast for foundational learning and that the analytical examples used in the materials assume cultural and procedural context that the candidate does not have. The result is a preparation experience that is structurally less effective for foreign-trained candidates than for ABA graduates, contributing to the pass rate gap documented in Paper No. 1.

Third, the post-admission integration period, beginning with bar examination passage and extending through the first several years of U.S. practice, lacks structured infrastructure calibrated to the foreign-trained lawyer's particular learning trajectory. Newly admitted foreign-trained lawyers entering large law firms typically have access to firm-internal training programs, mentorship arrangements, and a structured

caseload that supports learning by doing. Newly admitted foreign-trained lawyers entering solo, small-firm, in-house, or government practice, the modal entry pathway for foreign-trained lawyers in the United States, do not have the same access. The result is a bifurcated post-admission experience in which foreign-trained lawyers in well-resourced firm settings integrate efficiently and foreign-trained lawyers in less-resourced settings integrate slowly or not at all.

B. The bar examination performance gap

The bar examination performance gap is the most quantitatively visible feature of the integration problem. The data, drawn from the New York State Board of Law Examiners and the National Conference of Bar Examiners, has been stable across the most recent reporting periods. In February 2025, the foreign-educated pass rate on the New York bar examination was 30 percent against a domestic ABA graduate first-time pass rate of approximately 73 percent for out-of-state candidates. The full-year first-time pass rate for foreign-educated candidates in 2025 was approximately 45 percent, against approximately 80 percent for ABA graduates. The gap is consistent across administration years and across U.S. jurisdictions admitting foreign-trained lawyers.

The gap is the product of several compounding factors. The Uniform Bar Examination format, multistate multiple choice, multistate essay, and multistate performance test components, places significant weight on rapid issue spotting in legal areas where foreign-trained candidates have less prior exposure. The MPT component, which presents candidates with a closed-universe legal task to be completed in 90 minutes, draws heavily on U.S. legal-writing conventions and U.S. procedural framing that foreign-trained candidates may not have internalized. The MPRE, while less affected by foreign training profile, still tests U.S. professional responsibility doctrine that foreign-trained candidates have studied for less time than ABA graduates.

Standard commercial bar preparation programs do not, in their mass-market form, address the specific learning needs that drive the performance gap. They do not offer foundational instruction calibrated to foreign-trained candidate baselines; they do not offer comparative-method analytical bridges from foreign legal training to U.S. legal frameworks; they do not offer language and cultural-context support for candidates whose first language is not English; and they do not offer iterative, individualized feedback at the level of granularity that foreign-trained candidates need. AI-driven systems, on the analysis offered in Section IV, can address each of these calibration failures.

C. The post-admission integration gap

The post-admission integration gap is the least quantitatively documented of the three sub-problems but, in many respects, the most consequential for the operational capacity of the foreign-trained lawyer population. A foreign-trained lawyer who passes the U.S. bar examination has cleared a significant hurdle, but bar admission does not, by itself, produce a competent U.S. practitioner. Competent practice requires sustained learning across substantive law, procedural rules, professional conventions, client-development practices, and the operational tradecraft of legal services delivery. For foreign-trained lawyers entering small-firm, solo, or in-house settings, this learning is largely self-directed, with limited external support.

The post-admission gap manifests in several specific ways. In substantive law, foreign-trained lawyers entering U.S. practice frequently report uncertainty about the relative weight of statutory text, regulatory guidance, judicial precedent, and administrative practice in particular doctrinal areas, weights that vary across U.S. legal subfields and that are not always made explicit in the published materials. In procedural law, foreign-trained lawyers report a learning curve on the practical operation of state and federal procedural rules, particularly on the timing and form requirements that govern litigation practice. In professional conventions, foreign-trained lawyers report uncertainty about U.S. norms governing client communication, fee structures, conflict screening, and engagement letter drafting; the norms vary across firm settings, geographies, and practice areas, and are typically transmitted through mentorship rather than through formal documentation.

Each of these post-admission learning needs is, in principle, addressable through AI-driven systems calibrated to the foreign-trained lawyer's specific practice setting. A foreign-trained immigration lawyer in solo practice in Houston has different learning needs than a foreign-trained corporate lawyer in a small firm in New York; an AI system tailored to each profile can deliver context-specific learning content in a form that traditional CLE programming, addressed to broader audiences, cannot match. The policy framework set out in Section VI develops this proposition.

III. The Existing Framework

Sound policy proposals build on the existing legal and regulatory framework. This section sets out the principal U.S. authorities bearing on the use of AI in legal services and legal education, focused on the four most operationally relevant: ABA Formal Opinion 512, state bar guidance on generative AI, the NIST AI Risk Management Framework, and the federal workforce and immigration architecture within which foreign-trained lawyers operate.

A. ABA Formal Opinion 512 and the duties of competence and confidentiality

The American Bar Association Standing Committee on Ethics and Professional Responsibility issued Formal Opinion 512 on July 29, 2024, addressing the duties of lawyers using generative AI tools in the practice of law. The Opinion is the most authoritative U.S. professional responsibility pronouncement on the question and articulates the framework within which any AI-driven legal readiness system must operate.

The Opinion identifies six principal areas in which the existing Model Rules of Professional Conduct apply to lawyers' use of generative AI: the duty of competence under Rule 1.1, including the duty to maintain technological competence under Comment 8; the duty of confidentiality under Rule 1.6; the duties of communication and informed consent under Rule 1.4; the duties of supervision over lawyers and non-lawyers under Rules 5.1 and 5.3; the duty of candor toward the tribunal under Rule 3.3 and the related duty of meritorious claims under Rule 3.1; and the obligation of reasonable fees under Rule 1.5.

The Opinion's central holding is that lawyers using generative AI must understand the tools' capabilities and limitations, must protect client confidentiality including by carefully reviewing the data-handling practices of any AI provider, must obtain informed consent in appropriate cases, must verify the accuracy of AI-generated work product, and must supervise the AI tool's output as the lawyer would supervise the work of a junior associate or non-lawyer assistant. The Opinion does not prohibit lawyer use of generative AI; it conditions that use on the application of existing professional responsibility duties to the new technological context.

For purposes of the framework proposed in this paper, the Opinion is significant in two respects. First, the Opinion addresses generative AI in the practice of law, that is, in the delivery of legal services to clients. The proposed FTLA framework contemplates AI use in the integration and professional development of lawyers, that is, as a learning and credentialing tool rather than as a service-delivery tool. The two use cases share certain features but differ in others, and the FTLA framework is calibrated to the educational and integration context rather than to the service-delivery context. Second, the Opinion's emphasis on competence, confidentiality, and supervision is fully applicable to the educational context, and the proposed framework incorporates each of these requirements.

B. State bar guidance on generative AI

Beyond the ABA Opinion, multiple state bars have issued guidance on lawyer use of generative AI. The Florida Bar issued an ethics opinion in January 2024 addressing generative AI use in legal practice. The California State Bar issued practical guidance in November 2023. The New York State Bar Association

Task Force on Artificial Intelligence issued its Report and Recommendations in April 2024, addressing generative AI both in legal practice and in legal education. The Pennsylvania Bar Association, the District of Columbia Bar, and the State Bar of Texas have issued guidance addressing related questions.

The state guidance is broadly consistent with ABA Opinion 512 in its substantive direction. Common themes include the duty to understand the AI tool's operation, the duty to verify AI-generated content, the duty to maintain client confidentiality including by reviewing the AI provider's data-handling practices, and the obligation to disclose AI use to clients in appropriate circumstances. The state guidance varies in its specificity, with some states (notably Florida and California) offering more operationally detailed direction and others (including New York's Task Force Report) offering more analytical framing.

For the framework proposed in this paper, the state guidance is significant in establishing that the use of AI tools in legal contexts, properly conditioned, is consistent with state professional responsibility frameworks. No state has prohibited lawyer use of generative AI; the regulatory direction is uniformly toward conditional permission with disclosure, supervision, and verification requirements. A legal readiness system designed for foreign-trained lawyer integration, operating in the educational rather than service-delivery context, falls within the operational space that the state guidance contemplates.

C. The NIST AI Risk Management Framework

The National Institute of Standards and Technology released the AI Risk Management Framework (AI RMF 1.0) in January 2023, with the Generative AI Profile released as NIST AI 600-1 in July 2024. The Framework is voluntary, but it has become the de facto standard for AI risk management across U.S. federal agencies and a substantial share of private-sector AI deployments. The Framework articulates four functions, Govern, Map, Measure, and Manage, that organizations deploying AI systems should perform on an ongoing basis.

Of particular significance for legal-services and legal-education AI applications, the Framework addresses risks of bias, accuracy degradation, hallucination, and confidentiality compromise that are characteristic of large language model deployments. The Generative AI Profile (NIST AI 600-1) specifies risk management actions tailored to generative AI, including content provenance verification, harmful-output testing, and human oversight protocols. The Profile's recommendations are operationally detailed and provide a workable framework for the design of legal readiness systems.

The framework proposed in this paper incorporates the NIST AI RMF and the Generative AI Profile as the operational risk-management discipline for the FTLA Legal Readiness Platform. The four functions are mapped to specific design and deployment practices in Section VI.

D. Federal workforce and immigration architecture

Foreign-trained lawyers in the United States operate within a federal workforce and immigration framework that bears on their ability to access, complete, and apply U.S. legal training. Many foreign-trained lawyers enter the United States on F-1 student visas to pursue LL.M. study; some transition to H-1B specialty occupation visas, O-1 extraordinary ability visas, or employment-based permanent residence pathways for

post-admission practice; some hold L-1 intracompany transferee status as part of multinational law firm assignments. Each visa category imposes specific employment and educational restrictions that foreign-trained lawyer integration must navigate.

The Department of Labor administers the labor certification process applicable to certain employment-based immigration pathways. U.S. Citizenship and Immigration Services administers the visa adjudication process. The Department of State administers consular visa issuance for foreign-trained candidates entering the United States. The Department of Education's recognition of foreign credentials, while not directly determinative of bar admission, intersects with the foreign credentials evaluation that bar admission authorities conduct. Each of these agencies operates programs that, in their aggregate, shape the practical pathways through which foreign-trained legal expertise enters the U.S. profession.

The Department of Commerce, through SelectUSA and the International Trade Administration, has, as Paper No. 1 documented, treated inbound foreign direct investment as a national-interest objective for more than four decades. The legal-services workforce that supports inbound investment, encompassing both domestically trained and foreign-trained lawyers, is a component of the FDI infrastructure that SelectUSA has not, to date, formally engaged with. The framework proposed in this paper contemplates such engagement as part of its coordination component.

IV. Analysis

Sections II and III have established the integration problem and the legal and regulatory framework within which any AI-driven response must operate. This section analyzes the substantive case for AI deployment in foreign-trained lawyer integration, identifies the specific functions that AI is well suited to perform, identifies the functions that AI cannot perform and should not be assigned, addresses the equity and access dimension, and connects the analysis to the foreign direct investment objectives identified in FTLA Policy Paper No. 1.

A. Four legal-readiness functions AI can perform well

Four specific legal-readiness functions, on the analysis offered here, are well suited to AI-driven support and constitute the principal use cases for the proposed FTLA framework. Each is identified by reference to the integration sub-problems set out in Section II and to the comparative advantages of AI systems in performing personalized, comparative, asynchronous learning at scale.

1. Comparative-law translation

The first function is comparative-law translation: the identification, for a given U.S. doctrinal area, of the points of convergence and divergence with the foreign legal system in which the foreign-trained lawyer was originally trained. This function is at the heart of effective foreign-trained lawyer integration. A Brazilian-trained lawyer learning U.S. corporate law needs to understand how the U.S. system's fiduciary duty framework relates to the Brazilian Lei das Sociedades por Ações *dever de diligência* and *dever de lealdade*; a German-trained lawyer learning U.S. corporate governance needs to understand how the U.S. unitary board structure relates to the German two-tier system of *Aufsichtsrat* and *Vorstand* under codetermination requirements. The comparative analysis is, on the analysis offered in FTLA Policy Paper No. 1, the central cognitive activity of cross-border legal practice.

Generative AI systems trained on legal corpora can perform this comparative translation at a depth and at a scale that traditional educational infrastructure cannot match. A retrieval-augmented generation system with access to U.S. statutes and judicial decisions, foreign statutes in the major source jurisdictions of foreign-trained lawyers in the United States, and high-quality comparative-law treatises can produce, on demand, comparative analyses calibrated to the specific learner's foreign training background and the specific U.S. doctrinal area at issue. The output of such a system is not, of course, a substitute for the learner's own development of comparative-method competence; it is a learning tool that scaffolds the learner's development of that competence at a faster rate than would be possible without it.

The function is well suited to AI for three reasons. First, the underlying task is one of structured information retrieval and synthesis across well-defined corpora; this is the central capability of contemporary retrieval-augmented systems. Second, the function is asynchronous and personalized; the learner accesses the system at the time and in the depth that the learner's needs require, with the system adapting to the learner's prior interactions. Third, the function does not require the system to give legal advice to a client; it provides educational content to a learner who is, by hypothesis, already a trained lawyer (or a lawyer in training)

capable of evaluating the content critically. The use case sits squarely within the educational role for AI in legal contexts that ABA Opinion 512 contemplates.

2. Bar examination preparation calibrated to foreign-trained candidate profiles

The second function is bar examination preparation calibrated to foreign-trained candidate profiles. The structural mismatch between standard commercial bar preparation and foreign-trained candidate learning needs, identified in Section II.B, is, on the analysis offered here, addressable through AI-driven preparation systems that adapt to the candidate's specific foreign legal training, language profile, and identified learning gaps.

A calibrated AI bar preparation system would differ from standard commercial offerings in four respects. It would begin with a diagnostic assessment that maps the candidate's foreign legal training to the bar tested subjects, identifying the foundational gaps that the candidate needs to close. It would deliver foundational instruction in those gap areas, drawing on the U.S. common-law materials at the depth that the candidate's particular profile requires. It would offer iterative practice with bar tested question formats, including the multistate multiple choice, the multistate essay, and the multistate performance test, with feedback calibrated to the candidate's identified weaknesses. And it would adapt over time to the candidate's evolving performance, allocating more time to areas of continuing weakness and less to areas of demonstrated strength.

Each of these features is operationally feasible with current AI technology. The diagnostic assessment can be administered through a large language model with access to the candidate's foreign legal training summary and a calibrated assessment instrument. The foundational instruction can be delivered through retrieval-augmented generation drawing on standard bar preparation materials and supplementary sources. The iterative practice can be administered through prompts that present bar tested questions and evaluate candidate responses against scoring rubrics. The adaptive allocation can be managed through performance tracking and content sequencing logic. The combined system, deployed at scale, can deliver bar preparation at a level of personalization that no human-led program can cost-effectively provide.

The system does not, of course, take the bar examination on the candidate's behalf. The candidate's preparation, study, and examination performance remain the candidate's own. The AI system supports the preparation process; it does not substitute for it. This supplementary role is the role that the existing professional responsibility framework contemplates, and it preserves the substantive integrity of the bar examination as a measure of the candidate's preparedness for U.S. practice.

3. Regulatory-context onboarding for U.S. specialty practice

The third function is regulatory-context onboarding for U.S. specialty practice. The post-admission integration gap identified in Section II.C is, in significant part, a function of the time required for newly admitted foreign-trained lawyers to develop fluency in the regulatory architecture of their U.S. specialty practice. AI-driven systems can compress this development by delivering personalized, context-specific regulatory orientation calibrated to the practitioner's particular specialty and geography.

The function operates similarly to the comparative-law translation function but with a different focus. Where the comparative function maps U.S. doctrinal areas to foreign systems, the regulatory-context function maps U.S. doctrinal areas to U.S. operational practice: how the regulatory framework is administered, how agency guidance is interpreted, how the regulated community responds, what the recurring practical questions are, and how experienced practitioners typically address them. This operational dimension of legal practice is transmitted, in the traditional model, through mentorship and through in-firm training. AI systems can supplement, though not replace, this transmission, particularly for practitioners in solo, small-firm, in-house, or government settings where in-firm training is limited.

The function is particularly valuable in the cross-border specialty areas identified in FTLA Policy Paper No. 1: capital markets and securities, national security review (CFIUS), supply chain integrity (UFLPA), sanctions compliance (OFAC), anti-corruption compliance (FCPA), and cross-border corporate governance. Each of these specialties has a body of regulatory practice, agency guidance, enforcement precedent, and practitioner convention that the foreign-trained lawyer must develop fluency in. An AI system trained on the public regulatory record in each specialty can deliver personalized onboarding content that compresses the development timeline.

4. Ethics infrastructure for cross-border practice

The fourth function is ethics infrastructure for cross-border practice. FTLA Policy Paper No. 1 identified the development of a Cross-Border Practice Ethics Handbook as a programmatic priority. The Handbook addresses the recurring ethics questions in cross-border practice, the analysis of conflicts under both U.S. Model Rules and home-country bar rules, the management of confidentiality across jurisdictions with different privilege frameworks, the navigation of unauthorized-practice-of-law constraints for foreign-trained lawyers operating across state lines, and the handling of trust account and client funds in cross-border transactions.

AI-driven systems can support the deployment of the ethics infrastructure by delivering personalized ethics analysis on demand. A foreign-trained lawyer facing a recurring ethics question, a conflict between a U.S. corporate client and a home-country counterparty, for example, or a confidentiality question arising from cross-border discovery, can consult an AI system that has been trained on the relevant U.S. and home-country rules and can provide an initial analysis that the lawyer can then evaluate and act on. The AI system does not, of course, substitute for the lawyer's own ethical judgment or for consultation with bar counsel where the question is significant; it provides a structured analytical starting point that the lawyer can build on.

The function is particularly valuable for foreign-trained lawyers in solo and small-firm settings, who do not have access to the firm-internal ethics infrastructure that larger firms maintain. For these practitioners, an AI-driven ethics-support tool, used as a learning and analytical resource rather than as a substitute for professional judgment, can materially improve the quality of ethical analysis applied to recurring cross-border practice questions.

B. Four functions AI cannot perform and should not be assigned

The case for AI deployment in foreign-trained lawyer integration is bounded by clear limits on what AI systems can and should be asked to do. Four functions, in particular, fall outside the appropriate scope of AI deployment and should be reserved for human professionals.

First, AI systems should not give legal advice to clients on behalf of foreign-trained lawyers. The unauthorized-practice-of-law doctrine, the lawyer-client confidentiality framework, and the malpractice liability architecture all presuppose that legal advice is provided by an admitted lawyer who has assumed professional responsibility for the advice. AI systems do not satisfy these conditions and should not be used to give client-facing legal advice in any form. The proposed FTLA framework operates exclusively in the educational and professional development context and does not contemplate AI use in client-facing legal services.

Second, AI systems should not assess foreign-trained candidates' professional fitness for U.S. practice. Bar admission is a determination by a state's highest court, on the recommendation of the bar admission authority, that the candidate possesses the character, fitness, and competence required for U.S. practice. The determination involves judgment that an AI system is not equipped to make and that the existing professional responsibility framework does not authorize an AI system to make. The proposed FTLA framework supports candidate preparation; it does not assess candidate fitness, which remains a matter exclusively for human bar admission authorities.

Third, AI systems should not substitute for human mentorship in the foreign-trained lawyer integration process. Mentorship, identified in FTLA Policy Paper No. 1 as a core component of the FTLA programmatic response, depends on a sustained interpersonal relationship between mentor and mentee in which the mentor models professional conduct, transmits tacit knowledge that is not readily codified, and provides the kind of contextual judgment that AI systems cannot replicate. AI tools can supplement mentorship by providing learners with on-demand analytical resources between mentorship sessions, but they cannot substitute for the relational dimension of mentorship itself.

Fourth, AI systems should not be relied upon as authoritative sources of substantive law. The hallucination risk characteristic of large language models, the tendency to produce confident-sounding output that is factually incorrect or that cites non-existent authorities, is well documented and has been the subject of judicial sanctions in cases including *Mata v. Avianca, Inc.* (S.D.N.Y. 2023), where attorneys submitted briefs citing fabricated cases produced by ChatGPT. AI systems used in legal readiness applications must, in their design, mitigate this risk through retrieval-augmented architecture, citation verification, and explicit instruction to users that the AI output requires independent verification against primary authorities. The mitigation does not eliminate the risk but reduces it to a manageable level.

C. The equity and access dimension

The case for AI deployment in foreign-trained lawyer integration has an equity and access dimension that warrants explicit treatment. The traditional infrastructure for foreign-trained lawyer integration, LL.M. programs at leading U.S. law schools, commercial bar preparation programs, mentorship arrangements within well-resourced law firms, is characterized by significant cost and access barriers. LL.M. tuition at

the leading programs runs in excess of \$80,000 per academic year. Commercial bar preparation programs run in the range of \$3,000 to \$6,000 per administration. Mentorship within well-resourced firms is available primarily to candidates who secure positions in those firms, a population that is itself selected on prior advantages.

AI-driven legal readiness systems, deployed at scale, have the potential to improve access for foreign-trained lawyers who do not have the resources to access the traditional infrastructure at full price or in full form. The marginal cost of delivering AI-driven content to an additional learner is low; the system, once developed, can be made available to FTLA members at modest cost or, for certain components, at no cost. The result, on the analysis offered here, is an integration infrastructure that is more equitably accessible than the existing alternatives, particularly for foreign-trained lawyers from lower-income source jurisdictions or in circumstances of constrained personal finances.

The equity dimension is not, however, automatic. AI systems can also reproduce or amplify existing access inequities if they are designed and deployed without attention to access considerations. A system that requires high-bandwidth internet connectivity, recent computing hardware, or sophisticated user-interface fluency may be effectively inaccessible to learners who lack those resources. A system that has been trained on corpora that underrepresent certain source jurisdictions may produce comparative analysis that is more accurate for some learners than for others. A system that is offered exclusively in English may exclude learners whose English fluency is still developing. The framework proposed in this paper is calibrated to mitigate these access risks through design choices that are described in Section VI.

D. Economic stakes: the FDI legal services connection

The economic stakes of AI-driven legal readiness deployment connect to the foreign direct investment analysis developed in FTLA Policy Paper No. 1. That paper argued that foreign-trained lawyers integrated into U.S. practice are critical professional infrastructure for inbound foreign investment, and that the inefficiency of current integration pathways translates into deal friction and lost investment. The same logic applies, with refinements, to the AI-driven integration framework.

If AI-driven legal readiness systems compress the foreign-trained lawyer integration timeline, the resulting expansion of operational legal capacity translates into a more efficient cross-border legal services market and, by the analysis in Paper No. 1, into more efficient inbound foreign investment. The marginal effect on FDI volume is difficult to quantify precisely, but the directional analysis is clear: a foreign-trained lawyer who reaches operational fluency in U.S. cross-border practice in two years rather than four delivers, over a career, twice the cross-border legal capacity that the four-year-timeline lawyer delivers in the same period.

The aggregate effect across the foreign-trained lawyer population is, on a conservative estimate, equivalent to a substantial expansion of the U.S. cross-border legal services workforce without a corresponding expansion of the underlying population of foreign-trained lawyers. This expansion is, in operational terms, a productivity improvement, and productivity improvements in the legal-services workforce that supports inbound foreign investment are, by the framework articulated in FTLA Policy Paper No. 1, contributions

to U.S. capital markets integrity and to the realization of the inbound investment objectives that SelectUSA and analogous federal programs have pursued for more than four decades.

V. Comparative Perspective

The cross-border analytical method developed in FTLA Policy Paper No. 1 and applied in Paper No. 2 applies with full force to the question of AI in legal services and legal education. Several major jurisdictions have, in the past three years, articulated regulatory and professional frameworks for AI in legal contexts, and the U.S. framework should be read against these comparative developments. This section sketches the principal points of comparison.

A. The European Union: the AI Act and legal-services applications

The European Union's Artificial Intelligence Act (Regulation (EU) 2024/1689), addressed in FTLA Policy Paper No. 2 in connection with intellectual property questions, is also significant for legal-services and legal-education AI applications. The AI Act classifies AI systems by risk level and imposes obligations calibrated to that classification. The legal-services application of the Act turns on whether the particular system is classified as a high-risk AI system under Annex III, which includes systems used in education and vocational training and systems used in administration of justice.

Legal readiness systems used in foreign-trained lawyer integration may fall within the education and vocational training category if they are used for assessment of student performance or for determining access to educational institutions. Legal readiness systems used in lawyer professional development may fall outside the education category if they are used for self-directed learning by already-admitted lawyers. The classification analysis is fact-specific and depends on the particular features of the system. For systems classified as high-risk, the AI Act imposes obligations including risk management systems, data governance practices, technical documentation, transparency to users, human oversight, and accuracy and robustness requirements. The obligations are operationally significant and would, in practice, shape the design of any legal readiness system deployed in EU markets.

The Council of Bars and Law Societies of Europe (CCBE), the principal European bar association, has issued guidance on AI use by lawyers that is broadly aligned with ABA Opinion 512 in substantive direction but differs in operational specifics. The guidance addresses the duties of competence, confidentiality, and supervision in the European professional responsibility frameworks, with attention to the differences across European national bars in their treatment of these duties.

B. The United Kingdom: SRA guidance and the regulatory sandbox

The United Kingdom Solicitors Regulation Authority (SRA) has issued guidance on AI use by solicitors that is among the most operationally detailed of the major jurisdictions' guidance documents. The SRA's November 2023 thematic review of technology and innovation in the legal services sector, supplemented by subsequent risk outlook publications, addresses generative AI alongside other emerging technologies and articulates a framework for solicitor compliance. The SRA's regulatory sandbox, available since 2023, permits solicitors to test innovative AI-driven service models under regulatory supervision before full market deployment.

The UK approach is significant for two reasons. First, the SRA's guidance establishes a framework that is operationally compatible with both legal-services and legal-education AI applications, with the regulatory sandbox providing a mechanism for testing novel use cases including those that bear on professional development and integration. Second, the UK's broader post-Brexit positioning as a hospitable jurisdiction for AI innovation, reflected in the AI Regulation White Paper of March 2023 and the consultation on copyright and AI addressed in FTLA Policy Paper No. 2, creates an environment in which AI-driven legal readiness systems can be developed and tested with regulatory cooperation rather than regulatory friction.

The UK Bar Standards Board, which regulates barristers, has issued parallel guidance addressing AI use by barristers. The Bar Council, the representative body of the practicing bar, has produced detailed practical guidance for barristers on the use of generative AI in research, drafting, and practice management. The combined UK framework provides a workable model for the integration of AI tools into legal-services and legal-education infrastructure.

C. Singapore and Canada: targeted use cases

Beyond the EU and UK, several other jurisdictions have made notable contributions to the development of AI-driven legal-services and legal-education infrastructure. Singapore, through the Singapore Academy of Law and the Ministry of Law, has supported the development of GPT-Legal Q&A and related tools designed to support practitioner research and professional development. Singapore's broader positioning as a regional legal services hub has been complemented by deliberate investment in legal technology infrastructure, including the LawNet platform and adjacent tools.

Canada, through the Federation of Law Societies of Canada and provincial law societies, has issued guidance on AI use by lawyers that is broadly aligned with the ABA and SRA frameworks. The Law Society of Ontario, the largest provincial regulator, has been particularly active in publishing practical guidance on generative AI, including practice notes addressing competence, confidentiality, and disclosure obligations. Canada's federal jurisdiction includes the National Committee on Accreditation, which performs the foreign credentials evaluation that is the Canadian analog of the U.S. state bar foreign credentials evaluation under New York Rule 520.6 and similar rules.

The Singapore and Canada examples are significant for the framework proposed in this paper because they demonstrate the operational feasibility of AI-driven legal readiness systems deployed on terms compatible with traditional professional responsibility frameworks. The U.S. framework, while developed in a distinct regulatory environment, can draw on these comparative examples in its design and deployment choices.

VI. Policy Recommendations

The diagnosis offered in Sections II through V supports a policy framework with four components. Each addresses a specific integration sub-problem and is calibrated to operate within the existing professional responsibility, AI risk-management, and federal workforce frameworks. The framework is offered as a structural proposal for stakeholder consultation rather than as a fully specified deployment plan; detailed program documentation will be developed through the FTLA programmatic process.

A. The FTLA Legal Readiness Platform

FTLA proposes the development of an FTLA Legal Readiness Platform, an integrated AI-driven system supporting foreign-trained lawyer integration across the four functional areas identified in Section IV.A: comparative-law translation, bar examination preparation calibrated to foreign-trained candidate profiles, regulatory-context onboarding for U.S. specialty practice, and ethics infrastructure for cross-border practice. The Platform would be developed and operated by FTLA, with content and design contributions from FTLA members, affiliated U.S. law school faculty, and (where appropriate) external partners.

The Platform's design would incorporate the four functions of the NIST AI Risk Management Framework. Govern would establish the policies, accountability structures, and review processes governing Platform deployment. Map would identify the contexts of use, the user populations, and the risks associated with each component of the Platform. Measure would assess Platform performance against accuracy, fairness, transparency, and accessibility metrics on an ongoing basis. Manage would respond to identified risks through design adjustments, content corrections, and operational modifications.

The Platform would operate under explicit user-facing terms that establish its educational and professional-development role, that disclose its AI-driven nature, that identify the limits of its outputs, and that direct users to human professional resources for matters falling outside the Platform's appropriate scope. The terms would track the disclosure framework articulated in ABA Formal Opinion 512 and parallel state guidance, applied to the educational rather than service-delivery context.

The Platform would be available to FTLA Full Members and Associate Members at modest or no cost, with Affiliate and Institutional Members supporting the Platform's continued development through their membership contributions. The cost structure would, on the analysis offered in Section IV.C, address the equity and access dimension of foreign-trained lawyer integration by providing AI-driven readiness infrastructure on terms substantially more accessible than the traditional commercial alternatives.

B. AI-assisted bar preparation calibrated to foreign-trained candidates

Within the FTLA Legal Readiness Platform, a dedicated bar examination preparation module would address the structural mismatch between standard commercial bar preparation and foreign-trained candidate learning needs. The module would deliver the four features identified in Section IV.A.2: diagnostic assessment of foundational gaps, foundational instruction calibrated to candidate profile, iterative practice with bar tested question formats, and adaptive content sequencing in response to candidate performance.

The bar preparation module would operate as a complement to, not a substitute for, established commercial bar preparation programs. FTLA does not propose to replicate the full curriculum of major commercial providers; it proposes to address the specific calibration gap that those providers do not, in their mass-market form, fill. A foreign-trained candidate would, in the typical case, use the FTLA module alongside a commercial program, drawing on the FTLA module for foundational and comparative-method content and on the commercial program for the standard bar preparation curriculum.

The module would be developed in coordination with academic faculty at U.S. law schools whose LL.M. programs serve substantial foreign-trained populations. FTLA would seek formal liaison arrangements with the LL.M. program directors at the leading institutions and with the National Conference of Bar Examiners on the calibration of preparation content to bar examination architecture. The coordination would build on the academic coordination contemplated in FTLA Policy Paper No. 1, with the bar preparation module serving as a structured channel through which academic and bar admission stakeholders can engage with the foreign-trained candidate experience.

C. Cross-border practice ethics infrastructure

FTLA Policy Paper No. 1 identified the development of a Cross-Border Practice Ethics Handbook as a programmatic priority. This paper proposes that the Handbook be deployed not only as a static document but as a dynamic component of the FTLA Legal Readiness Platform, with AI-driven analytical support that delivers personalized ethics analysis on demand.

The dynamic ethics infrastructure would operate as follows. A foreign-trained lawyer facing a recurring cross-border ethics question would consult the Platform, providing the relevant facts (suitably anonymized to protect client confidentiality). The Platform would deliver an initial analysis drawing on the relevant U.S. Model Rules, the corresponding home-country bar rules in the principal source jurisdictions, the CCBE Charter of Core Principles, the IBA International Principles on Conduct for the Legal Profession, and the leading scholarship on cross-border practice ethics. The analysis would be presented as an analytical starting point that the lawyer evaluates and acts on, with the Platform directing the lawyer to bar counsel consultation, to FTLA Affiliate Members in the relevant practice area, or to other appropriate human resources for significant questions.

The ethics infrastructure would, by design, operate on the educational rather than service-delivery side of the line that ABA Opinion 512 articulates. The Platform does not provide ethics advice to the lawyer's clients; it provides analytical support to the lawyer in the lawyer's own ethical analysis. The lawyer's professional responsibility for the resulting decisions remains the lawyer's own. This structural framing ensures that the dynamic ethics infrastructure operates within the existing professional responsibility framework rather than asking that framework to expand to accommodate a new use case.

D. Coordination with federal workforce and FDI agencies

FTLA Policy Paper No. 1 contemplated coordination with federal economic development apparatus, including the Department of Commerce's SelectUSA program, on the legal-services dimension of inbound

foreign investment. This paper extends that contemplation to the AI-driven legal readiness infrastructure proposed here. The coordination would operate through three principal channels.

First, FTLA would seek liaison arrangements with SelectUSA and the International Trade Administration on the role of AI-driven legal readiness in the broader inbound foreign investment infrastructure. The argument, developed in Section IV.D, is that productivity improvements in the foreign-trained lawyer integration process translate into more efficient cross-border legal services and, by the analysis in FTLA Policy Paper No. 1, into more efficient inbound foreign investment. Federal economic development apparatus has, on this view, a direct interest in the success of AI-driven legal readiness deployment.

Second, FTLA would seek liaison arrangements with the Department of Labor on the workforce dimension of foreign-trained lawyer integration. AI-driven legal readiness systems, by compressing the integration timeline, expand the operational legal-services workforce in the cross-border practice areas of greatest national-interest priority. The Department of Labor's role in workforce development and labor market analysis intersects with the FTLA framework in ways that warrant structured engagement.

Third, FTLA would seek liaison arrangements with the National Institute of Standards and Technology on the application of the AI Risk Management Framework to legal-services and legal-education AI applications. NIST has, in its Generative AI Profile and adjacent publications, indicated continuing engagement with sector-specific AI risk-management questions. The legal-services and legal-education context presents a discrete set of risk-management questions that NIST engagement could helpfully address.

VII. Implementation Considerations

The four policy recommendations set out in Section VI are designed for implementation through existing institutional channels, without the creation of new federal agencies or comprehensive regulatory frameworks. This section addresses the principal implementation considerations: institutional capacity, technical architecture, governance arrangements, transition planning, and risk.

Institutional capacity. The FTLA Legal Readiness Platform, the central institutional vehicle of the proposed framework, would be developed and operated by FTLA in coordination with the academic, bar admission, and federal partners identified in Section VI. The institutional capacity required is substantial but not unprecedented; comparable platforms have been developed by professional associations in adjacent fields, including the American Institute of Certified Public Accountants and the American Medical Association in their respective continuing professional development infrastructures. FTLA would build the Platform incrementally, beginning with the comparative-law translation and bar preparation modules in 2026 and expanding to the regulatory-context and ethics modules in 2027 and 2028.

Technical architecture. The Platform's technical architecture would prioritize retrieval-augmented generation over pure large-language-model output, with the AI components drawing on curated corpora rather than producing content from model weights alone. The retrieval-augmented architecture mitigates the hallucination risk identified in Section IV.B and ensures that Platform output is anchored in identifiable primary sources that the user can verify independently. The architecture would also incorporate citation verification at the output stage, with the Platform validating that cited authorities exist and are accurately characterized before delivering output to the user. The combined architecture is more resource-intensive than pure language-model deployment but is, on the analysis offered here, the appropriate architecture for legal-services and legal-education applications.

Governance arrangements. The Platform's governance would follow the four functions of the NIST AIRMF as set out in Section VI.A. A FTLA AI Governance Committee, composed of FTLA leadership, affiliated faculty, practitioner representatives, and outside ethics counsel, would oversee Platform development and deployment. The Committee would publish annual reports on Platform performance, identified risks, and corrective actions taken. The reports would be available to FTLA members and (in suitably aggregated form) to the broader U.S. legal profession and policy community.

Transition planning. The Platform's deployment would be staged to allow for iterative refinement based on user experience and identified issues. The initial release would be designated as a beta deployment with explicit user-facing terms acknowledging the developmental status and inviting user feedback. After a defined beta period, the Platform would transition to general availability for FTLA members. Throughout the transition, FTLA would maintain capacity to suspend or modify Platform components in response to identified issues, including issues identified through the bar admission and professional responsibility frameworks within which the Platform operates.

Risks. The principal risks of the proposed framework are five. First, the hallucination risk, that Platform output may include factually incorrect content, particularly fabricated case citations or doctrinal

misstatements, is real and must be managed through the technical architecture and user-facing disclosure described above. Second, the access risk, that the Platform may, despite its design intent, prove less accessible to certain learner populations than to others, must be managed through deliberate design choices addressing language, technical accessibility, and cost barriers. Third, the substitution risk, that Platform users may rely on the Platform in lieu of human professional consultation in matters where human consultation is essential, must be managed through user-facing framing and through directed referral to human resources for significant questions. Fourth, the regulatory risk, that the Platform may, in particular jurisdictions or applications, be characterized as engaged in the unauthorized practice of law or as otherwise running afoul of professional responsibility frameworks, must be managed through ongoing engagement with bar regulators and through the educational rather than service-delivery framing of Platform operation. Fifth, the reputational risk, that errors or controversies in Platform operation may adversely affect FTLA's broader institutional standing, must be managed through the governance arrangements and transparency practices described above.

These risks are real but manageable. They do not, in the analysis offered here, outweigh the benefits of the framework: a more efficient foreign-trained lawyer integration process, a more accessible professional development infrastructure for the foreign-trained lawyer population, and a more productive cross-border legal-services workforce serving the U.S. national-interest objectives identified in FTLA Policy Paper No. 1. The framework is offered as a structured response to a workforce-integration problem that the United States has not, to date, addressed at the level of operational specificity that the problem demands.

VIII. Conclusion

This paper has argued that artificial intelligence, deployed as workforce-integration infrastructure rather than as a substitute for legal labor, can address several of the specific pressure points identified in FTLA Policy Paper No. 1 in the integration of foreign-trained lawyers into the U.S. legal profession. The paper has identified the integration sub-problems that AI is well suited to address (comparative-law translation, calibrated bar preparation, regulatory-context onboarding, and ethics infrastructure for cross-border practice), the limits beyond which AI deployment should not extend, and the policy framework through which FTLA, in coordination with academic, bar admission, and federal partners, can deploy AI-driven legal readiness infrastructure on terms consistent with the existing professional responsibility, AI risk-management, and federal workforce frameworks.

Three points warrant emphasis as the paper closes.

First, the framing offered here is consistent with the established U.S. policy posture on AI in legal services. The American Bar Association's Formal Opinion 512, the parallel state bar guidance, and the NIST AI Risk Management Framework collectively articulate a regulatory direction in which AI use in legal contexts is conditionally permitted on terms calibrated to the existing duties of competence, confidentiality, and supervision. The framework offered in this paper operates within that regulatory direction; it does not require any expansion of the existing permissions or any modification of the existing duties. The framework is, in the precise sense, an institutional application of an existing regulatory framework rather than a regulatory innovation.

Second, the framing offered here is grounded in the U.S. national interest in a functioning cross-border legal-services workforce, not in the interests of foreign-trained lawyers as such. The case for the framework is not that foreign-trained lawyers deserve more efficient integration infrastructure because of who they are; it is that the United States gets better outcomes, more efficient cross-border legal services, more reliable inbound foreign investment processes, more rigorous compliance with the federal cross-border architecture, when the integration of foreign-trained lawyers into U.S. practice is conducted with the technological infrastructure that contemporary AI systems make available. The case is in the same register as the inbound investment case developed in FTLA Policy Paper No. 1 and the intellectual property case developed in FTLA Policy Paper No. 2: the United States has, in foreign-trained lawyers integrated into U.S. practice, an asset whose value the United States has not yet learned to use as fully as it should.

Third, the work proposed here is institutional and incremental, not regulatory in any expansive sense. The policy framework set out in Section VI does not call for new federal agencies, new statutory frameworks, or comprehensive regulation of AI in legal services. It calls for the development of a professional-association-operated platform deployed within the existing professional responsibility framework, supported by coordination arrangements with existing federal economic-development and workforce apparatus. The marginal cost of implementing this framework is modest. The marginal benefit, by the analysis offered, is substantial: a foreign-trained lawyer integration process that compresses an integration timeline currently measured in years to one measured in months, with cumulative benefits to U.S. capital

markets, foreign investment, and cross-border compliance integrity that compound over the careers of the integrated practitioners.

The Foreign Trained Lawyers Association will, over the months and years following the publication of this paper, develop the FTLA Legal Readiness Platform in stages, beginning with the comparative-law translation and bar preparation modules in 2026. The development will be conducted in coordination with the academic, bar admission, and federal partners identified in Section VI, and the resulting Platform will be made available to FTLA members on the access terms identified in Section VI.A. The Association looks forward to engaging with the broader U.S. legal profession, the bar admission community, federal economic-development apparatus, and the AI policy community on the framework offered in this paper, and to contributing the practitioner-grounded perspective that the foreign-trained lawyer population brings to the rapidly evolving field of AI in legal services.

References

This paper draws principally on U.S. government sources, professional association publications, and authoritative comparative material from the European Union, the United Kingdom, Canada, and Singapore. Citations to bar examination data are drawn from the New York State Board of Law Examiners and the National Conference of Bar Examiners. URLs and pinpoint citations are accurate as of the date of publication and may evolve thereafter.

1. American Bar Association Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512: Generative Artificial Intelligence Tools (July 29, 2024).
2. American Bar Association, Model Rules of Professional Conduct (as amended through 2024).
3. Florida Bar, Ethics Opinion 24-1 (Jan. 19, 2024).
4. California State Bar Standing Committee on Professional Responsibility and Conduct, Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law (Nov. 2023).
5. New York State Bar Association Task Force on Artificial Intelligence, Report and Recommendations (April 2024).
6. District of Columbia Bar, Ethics Opinion on Generative AI (2024).
7. Pennsylvania Bar Association, Joint Formal Opinion 2024-200, Ethical Issues Regarding the Use of Artificial Intelligence (May 22, 2024).
8. State Bar of Texas, Practice Guidance on Artificial Intelligence (2024).
9. National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (January 2023).
10. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (July 2024).
11. *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023).
12. Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), 2024 O.J. (L 168) 1.
13. Council of Bars and Law Societies of Europe (CCBE), Guidance on the Use of Artificial Intelligence by Lawyers (2024).
14. Solicitors Regulation Authority (UK), Risk Outlook Report: The Use of Artificial Intelligence in the Legal Market (Nov. 2023).
15. Bar Standards Board (UK), Guidance on the Use of ChatGPT and other Generative AI Software (2024).
16. Federation of Law Societies of Canada, Guidance on the Use of Artificial Intelligence in the Practice of Law (2024).
17. Law Society of Ontario, Practice Guideline on the Use of Artificial Intelligence (2024).
18. Singapore Academy of Law, Guide to Harnessing Generative AI for the Legal Profession (2023).
19. New York State Board of Law Examiners, Press Release, February 2025 Examination Results.
20. New York State Board of Law Examiners, NY Bar Exam 2025 Statistics.

21. Rules of the Court of Appeals for the Admission of Attorneys and Counselors at Law, 22 N.Y.C.R.R. § 520.6.
22. National Conference of Bar Examiners, Comprehensive Guide to Bar Admission Requirements (2025 ed.).
23. U.S. Department of Commerce, SelectUSA Investment Summit Materials (2025-2026).
24. Executive Order 13859, Maintaining American Leadership in Artificial Intelligence (Feb. 11, 2019).
25. Executive Order 14110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (Oct. 30, 2023).
26. Mariami Zhorzholiani, Authorship, Originality, and the Machine-Generated Work: U.S. Copyright Doctrine, Patent Disclosure, and the Governance of Generative Artificial Intelligence, FTLA Policy Paper No. 2 (April 2026).
27. Foreign Trained Lawyers Association, Policy Paper No. 1: Foreign Trained Legal Expertise as Critical Infrastructure for U.S. Cross-Border Compliance, Capital Markets Integrity, and Inbound Foreign Direct Investment (April 2026).

Suggested citation: Sopiko Gugava, Artificial Intelligence as Legal Workforce Infrastructure: Integrating Foreign-Trained Lawyers into the U.S. Economy through AI-Driven Legal Readiness Systems, FTLA Policy Paper No. 3 (Apr. 2026).

© 2026 Foreign Trained Lawyers Association. All rights reserved. This paper may be reproduced and distributed for noncommercial purposes with attribution to the Association and the author.

