AIR GAP VOICE

Attorney-Client Privilege and AI Dictation

April 2026

PUBLIC RELEASE

This paper examines whether the architecture of an AI dictation tool, specifically, whether it transmits audio to external servers or processes it on the attorney's own device, carries consequences for attorney-client privilege, the work product doctrine, and the attorney's ethical obligations under the Model Rules of Professional Conduct. It concludes that local processing eliminates, rather than merely mitigates, the privilege risks inherent in cloud-based dictation.

Author: Marvin F. L. Hansen
Contact: marvin.hansen@airgapvoice.com
Web: www.airgapvoice.com

Table of Contents

© 2026 Air Gap Voice | Public Release

1. Introduction

Artificial intelligence has arrived in the practice of law, and its most immediate, practical application is dictation. Attorneys have always dictated: to secretaries, into tape recorders, through digital voice recorders. The modern iteration of this centuries-old practice routes spoken words through AI-powered speech recognition engines that convert voice to text with remarkable accuracy. The productivity gains are substantial and well-documented.

The question this paper addresses is not whether attorneys should use AI dictation. They already do, and the efficiency arguments are compelling. The question is whether the particular architecture of a dictation tool — specifically, whether it transmits audio to external servers or processes it entirely on the attorney's own device — carries consequences for attorney-client privilege, the work product doctrine, and the attorney's ethical obligations under the Model Rules of Professional Conduct.

The answer, as this paper will demonstrate, is that it does. The distinction between cloud-based and locally processed dictation is not a matter of technical preference. It is a question of whether the attorney has voluntarily disclosed privileged communications to a third party, and whether that disclosure constitutes a waiver of privilege that no encryption protocol or vendor assurance can fully remedy.

This analysis proceeds from established privilege doctrine, applies it to the specific factual circumstances of cloud-based dictation, and examines the ethical framework governing an attorney's technology choices. It concludes with an assessment of local processing as the architecture that eliminates, rather than merely mitigates, the privilege risks inherent in transmitting spoken legal work to third-party servers.

2. The Privilege Framework

2.1 Attorney-Client Privilege Fundamentals

The attorney-client privilege is the oldest recognised privilege for confidential communications in the common law. Its purpose, as the Supreme Court articulated in Upjohn Co. v. United States, 449 U.S. 383 (1981), is to encourage "full and frank communication between attorneys and their clients and thereby promote broader public interests in the observance of law and administration of justice." The privilege protects confidential communications made for the purpose of obtaining or providing legal advice.

The elements are well-settled. The communication must be made between privileged persons (attorney and client), in confidence, for the purpose of obtaining or providing legal assistance. The privilege attaches only to communications made with a reasonable expectation of confidentiality. When that expectation is defeated — whether by the presence of third parties, voluntary disclosure, or the failure to take reasonable precautions — the privilege may be waived.

2.2 The Work Product Doctrine

Distinct from but related to the attorney-client privilege is the work product doctrine, established in Hickman v. Taylor, 329 U.S. 495 (1947). The doctrine protects materials prepared by an attorney in anticipation of litigation or for trial. This includes not only documents but also the attorney's mental impressions, conclusions, opinions, and legal theories — collectively termed "opinion work product" — which receives near-absolute protection.

When an attorney dictates case strategy, litigation analysis, deposition preparation notes, or settlement evaluations, the resulting text constitutes work product. If the dictation captures the

© 2026 Air Gap Voice | Public Release

attorney's mental impressions and legal theories, it is opinion work product entitled to the highest level of protection. The relevance to dictation technology is direct: the spoken words that become written work product pass through whatever infrastructure the dictation tool employs. If that infrastructure includes third-party servers, the work product has been transmitted to a third party.

2.3 The Reasonable Expectation of Confidentiality

The privilege requires that the communication be made with a reasonable expectation of confidentiality. This standard is objective, not subjective. It is not enough that the attorney believed the communication was confidential; the circumstances must be such that a reasonable person would have expected confidentiality. Courts evaluate the totality of the circumstances, including the steps taken to maintain confidentiality and the nature of the communication channel employed.

This standard creates a direct intersection with technology choices. An attorney who dictates privileged information into a system that transmits that information to external servers must demonstrate that the transmission channel supports a reasonable expectation of confidentiality. The burden of establishing that expectation falls on the party asserting the privilege.

2.4 Voluntary Disclosure as Waiver

The most consequential principle for this analysis is the doctrine of voluntary disclosure. When privileged information is voluntarily disclosed to a third party who is not within the privilege, the privilege is waived. The disclosure need not be intentional in the sense that the disclosing party meant to waive the privilege; it is sufficient that the disclosure was voluntary and that it resulted in the information being communicated to a non-privileged party.

Federal Rule of Evidence 502 provides some protection against inadvertent disclosure, but its protections are limited and apply primarily to disclosure made in the context of federal proceedings. The rule does not protect against systematic, knowing transmission of privileged content to third-party service providers as a routine feature of the attorney's workflow. When an attorney elects to use a cloud dictation service, the transmission of audio to that service's servers is not inadvertent; it is a designed and expected feature of the tool.

3. Cloud Dictation as Potential Waiver

3.1 The Transmission Problem

When an attorney dictates into a cloud-based speech recognition service, the following occurs: the attorney's spoken words — which may include privileged communications, work product, or both — are captured by a microphone, converted to a digital audio stream, and transmitted over a network to servers operated by the dictation vendor. On those servers, the audio is processed by the vendor's speech recognition engine, converted to text, and the text is returned to the attorney's device.

During this process, a complete audio recording of the attorney's dictation exists on infrastructure controlled by the vendor. The vendor's employees, systems, and subprocessors have potential access to that audio. The duration of retention, the purposes for which the audio may be used, and the circumstances under which it may be disclosed are governed not by the attorney's professional obligations but by the vendor's terms of service and privacy policy.

This is, in substance, a voluntary transmission of privileged communications to a third party. The attorney chose to use the tool. The attorney chose to speak privileged information into the tool.

© 2026 Air Gap Voice | Public Release

The tool, by design, transmitted that information to servers outside the attorney's control. The transmission was not accidental; it was the mechanism by which the tool functions.

3.2 The Third-Party Doctrine

The third-party doctrine, while primarily a Fourth Amendment concept from Smith v. Maryland, 442 U.S. 735 (1979), and United States v. Miller, 425 U.S. 435 (1976), has conceptual relevance to the privilege analysis. Under this doctrine, information voluntarily conveyed to a third party carries a diminished expectation of privacy because the individual has assumed the risk that the third party may disclose it.

An opposing party challenging privilege could argue that the attorney's use of cloud dictation constituted voluntary disclosure to the vendor, defeating the reasonable expectation of confidentiality. Whether a court would accept this argument depends on the specific facts, but the argument is not frivolous, and the attorney who must defend against it bears the burden of proof.

3.3 Vendor Terms of Service

An attorney using a cloud dictation service is bound by the vendor's terms of service. These terms routinely include provisions that should give any practising attorney pause. Common provisions include broad licences to use submitted data for service improvement, the right to use aggregated or anonymised data without restriction, retention of data for unspecified periods, and the right to disclose data in response to legal process or government requests.

The attorney who accepts these terms has agreed, contractually, that the vendor may use the dictated content for purposes beyond transcription. And the attorney has agreed that the vendor may disclose the content in response to a subpoena or court order — a disclosure that the attorney would have no standing to challenge, because the records belong to the vendor, not to the attorney.

3.4 The Reasonable Precautions Standard

Model Rule 1.6(c) requires attorneys to "make reasonable efforts to prevent the inadvertent or unauthorised disclosure of, or unauthorised access to, information relating to the representation of a client." ABA Formal Opinion 477R (2017) elaborated on this standard, identifying factors relevant to determining whether an attorney's efforts are reasonable, including the sensitivity of the information, the likelihood of disclosure if additional safeguards are not employed, and the difficulty of implementing the safeguards.

This is a critical point. The reasonableness inquiry is not conducted in a vacuum. When a readily available alternative eliminates the risk that the precaution is meant to address, the failure to adopt that alternative becomes relevant to whether the attorney's efforts were, in fact, reasonable.

3.5 State Bar Ethics Opinions

Multiple state bar associations have addressed the ethical implications of cloud computing and third-party technology services for attorneys. The New York State Bar Association (Opinion 842, 2010), the State Bar of California (Formal Opinion 2010-179), and the Florida Bar Board of Governors (Opinion 12-3, 2013) have each concluded that attorneys may use cloud services, provided they take reasonable care to ensure confidentiality is maintained and understand the vendor's data handling practices, security measures, and policies regarding data access and disclosure.

Each of these opinions was issued before AI dictation became prevalent, and none squarely addresses the specific risk profile of streaming live audio of privileged attorney communications to third-party servers for real-time processing. The risk profile of dictation is arguably more acute

© 2026 Air Gap Voice | Public Release

than the risk profile of document storage, because dictation involves the real-time transmission of the attorney's unfiltered spoken thoughts, including the kind of candid, off-the-cuff analysis that constitutes the most sensitive category of work product.

3.6 Encryption in Transit Versus Elimination of Transit

Proponents of cloud dictation services frequently cite encryption as a sufficient safeguard. The audio is encrypted in transit, encrypted at rest on the vendor's servers, and accessible only to authorised personnel. This argument conflates mitigation with elimination.

Encryption in transit protects against interception by unauthorised third parties during transmission. It does not address the fundamental issue: the data arrives at a destination outside the attorney's control. At the vendor's servers, the audio must be decrypted for processing; speech recognition cannot operate on encrypted audio. During processing, the content exists in an unencrypted state on the vendor's infrastructure.

> **The distinction between encrypting a transmission and eliminating the transmission is not semantic. Encryption mitigates the risk of interception. Elimination of transmission eliminates the risk of interception, the risk of vendor access, the risk of vendor disclosure, the risk of subpoena to the vendor, and the risk of breach at the vendor. These are categorically different risk profiles.**

4. The Metadata Problem

Even if one accepts the argument that encrypted transmission of dictation audio adequately protects the content of privileged communications, a separate and significant problem remains: metadata.

4.1 What Metadata Cloud Dictation Generates

Cloud dictation services necessarily generate metadata in the course of their operation. This metadata may include timestamps indicating when dictation sessions occurred and their duration, the frequency and pattern of dictation usage, device identifiers, network information including IP addresses, language and vocabulary settings, and in some cases, audio characteristics sufficient to identify the speaker.

This metadata is generated on the vendor's infrastructure and is within the vendor's possession and control. It is typically not covered by attorney-client privilege, because it does not constitute a confidential communication between attorney and client; it is a business record of the vendor documenting the attorney's use of the vendor's service.

4.2 Metadata as Discoverable Information

Metadata held by a third-party vendor is subject to discovery. A party in litigation could subpoena the dictation vendor for records of the opposing attorney's dictation activity. The resulting production could reveal when the attorney worked on the case, for how long, with what frequency, and potentially from what location. This information, while not itself privileged content, can provide significant strategic intelligence.

4.3 Pattern-of-Life Analysis

The aggregation of dictation metadata over time enables what intelligence professionals term "pattern-of-life analysis." Even without access to the content of the dictation, the metadata reveals behavioural patterns. A surge in dictation activity may indicate intensive case preparation. A shift

in dictation timing may indicate a change in strategy or staffing. The commencement of dictation activity on a new matter may reveal the existence of a previously undisclosed engagement.

As Chief Justice Roberts observed in Carpenter v. United States, 585 U.S. 296 (2018), metadata can provide "an intimate window into a person's life, revealing not only his particular movements, but through them his familial, political, professional, religious, and sexual associations." While this observation was made in the Fourth Amendment context, its logic applies with equal force to the confidentiality concerns of legal practice.

> **A dictation tool that processes locally generates no metadata on external servers. There are no vendor records to subpoena, no timestamps to analyse, no usage patterns to reconstruct. The metadata problem is not mitigated; it is absent.**

5. Ethical Obligations

5.1 Model Rule 1.6: Confidentiality of Information

Model Rule 1.6(a) provides that a lawyer "shall not reveal information relating to the representation of a client unless the client gives informed consent, the disclosure is impliedly authorised in order to carry out the representation, or the disclosure is permitted by [enumerated exceptions]." The rule's scope is broad: it protects all information relating to the representation, regardless of its source.

Model Rule 1.6(c), added in 2012, imposes an affirmative obligation: a lawyer "shall make reasonable efforts to prevent the inadvertent or unauthorised disclosure of, or unauthorised access to, information relating to the representation of a client." This provision transformed the duty of confidentiality from a purely negative obligation (do not disclose) to a positive one (take reasonable steps to prevent disclosure).

Applied to dictation technology, this analysis favours local processing: the information is highly sensitive (attorney's candid legal analysis), the likelihood of disclosure through cloud processing is non-trivial, and the alternative (local processing) is readily available and imposes no meaningful limitation on the attorney's ability to represent clients.

5.2 Model Rule 1.1 and the Duty of Technological Competence

Model Rule 1.1 requires that a lawyer provide competent representation, which includes "the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation." Comment 8 to the rule, amended in 2012, specifies that competent representation requires the lawyer to "keep abreast of changes in the law and its practice, including the benefits and risks associated with relevant technology."

ABA Formal Opinion 483 (2018) addressed the duty of technological competence in the context of data breaches, emphasising that attorneys must understand the technology they use and take reasonable steps to minimise the risk of unauthorised disclosure. Applied to dictation, the duty of technological competence requires the attorney to understand what happens to dictated audio. An attorney who uses cloud dictation without understanding its architecture, or without evaluating whether a local alternative is available, may fail to meet this standard.

© 2026 Air Gap Voice | Public Release

5.3 The Reasonable Efforts Standard

ABA Formal Opinion 477R (2017) provided extensive guidance on what constitutes "reasonable efforts" to protect client information in electronic communications. The opinion identified a non-exhaustive list of factors, including the nature of the threat, how client confidential information is transmitted and stored, the use of reasonable electronic security measures, and whether the attorney has the ability to assess the level of security employed.

In the dictation context, the reasonable efforts analysis is straightforward. The information at issue — the attorney's spoken words during legal work — is among the most sensitive categories of information in legal practice. The risk of cloud processing is identifiable and well-understood. And an alternative exists that eliminates the risk entirely. Under these circumstances, the failure to evaluate and consider that alternative is difficult to characterise as a reasonable effort.

5.4 Supervisory Obligations: Rules 5.1 and 5.3

Model Rule 5.1 requires partners and supervising attorneys to "make reasonable efforts to ensure that the firm has in effect measures giving reasonable assurance that all lawyers in the firm conform to the Rules of Professional Conduct." Model Rule 5.3 extends similar obligations with respect to nonlawyer assistants, including technology vendors whose services the firm employs.

For law firm managing partners, chief information officers, and legal technology committees, these rules create an obligation to evaluate the firm's dictation technology from a privilege and confidentiality perspective. If the firm deploys a cloud-based dictation solution across its practice groups, the firm's leadership bears responsibility for ensuring that the deployment is consistent with the attorneys' ethical obligations.

6. The Case for Local Processing

The analysis to this point has identified a series of privilege, work product, and ethical risks associated with cloud-based dictation. Each of these risks shares a common origin: the transmission of the attorney's spoken words to infrastructure outside the attorney's control. Local processing — dictation technology that converts speech to text entirely on the attorney's own device, with no network transmission of any kind — eliminates that common origin.

6.1 Eliminating the Waiver Question

When dictation is processed locally, no audio leaves the attorney's device. There is no transmission to a third party. There is no voluntary disclosure. The waiver analysis that dominates the discussion of cloud dictation simply does not arise. This is not a mitigation of waiver risk. It is the elimination of the factual predicate on which a waiver argument depends.

6.2 No Vendor Terms to Parse

A locally processed dictation tool that requires no account creation, no login, and no network connection presents no terms-of-service issue. There is no vendor agreement granting rights to use dictation data. There is no privacy policy permitting data retention. The attorney's dictation exists only on the attorney's device, subject only to the attorney's own data management policies and the firm's information governance framework.

© 2026 Air Gap Voice | Public Release

6.3 No Subpoena Risk to Vendor

If dictation data never reaches a vendor's servers, there are no vendor-held records to subpoena. The metadata problem identified in Section 4 is absent. No third party possesses records of when the attorney dictated, for how long, or with what frequency. The attorney's workflow remains within the attorney's control and within the protections of privilege and work product doctrine.

6.4 Verifiable Compliance

A particularly significant advantage of local processing is its verifiability. An attorney relying on a cloud vendor's assurances of security is, ultimately, taking the vendor at its word. The attorney cannot independently verify what happens to audio once it reaches the vendor's servers.

With local processing, the attorney can verify the behaviour independently. Network monitoring tools, available on every modern operating system, can confirm that the dictation application makes no network connections during operation. The attorney can operate the tool in airplane mode or on a machine with no network interface and confirm that it functions identically. This verifiability transforms the attorney's security posture from trust-based to evidence-based.

6.5 The Practical Dimension

Beyond the legal analysis, local processing offers a practical benefit that directly serves the purpose underlying the privilege: encouraging full and frank communication. An attorney who dictates into a cloud service carries a cognitive overhead — however slight — of wondering whether the transmission is truly secure. An attorney who knows that the dictation is processed locally, with no network transmission, can dictate freely, including the kind of candid, unfiltered analysis that is most valuable in legal work and most in need of protection.

7. Implementation Considerations for Law Firms

7.1 Firm-Wide Deployment

A locally processed dictation solution that requires no cloud infrastructure, no user accounts, and no network connectivity can be deployed through the firm's standard software distribution mechanisms. There are no servers to provision, no APIs to configure, and no vendor integrations to maintain. The firm's IT department installs the application on each attorney's device, and the deployment is complete.

7.2 No Credentials to Manage

Cloud dictation services typically require individual user accounts, which create their own security considerations: credential management, password policies, multi-factor authentication, and the risk of unauthorised access through compromised credentials. A local dictation tool that requires no accounts eliminates this entire category of administrative and security overhead. There are no credentials to provision, rotate, or deactivate when personnel depart.

7.3 Integration with Document Workflows

Attorneys require that dictated text appear in the document they are working on, at the cursor position, in the application they are using. Solutions that achieve this by inserting text directly at the system input focus — rather than routing through a clipboard or intermediate application — integrate seamlessly with existing workflows. The attorney dictates, the text appears in the document. The experience is indistinguishable from cloud dictation in terms of usability, with none of the privilege implications.

© 2026 Air Gap Voice | Public Release

7.4 Training Considerations

The learning curve for local dictation tools that employ modern speech recognition models is minimal. Attorneys who have used any dictation tool will find the experience familiar. The principal training point is not operational but conceptual: ensuring attorneys understand why the firm has selected a locally processed solution, and what that choice means for their privilege and confidentiality posture.

7.5 Compliance Documentation

For firms that must demonstrate their technology choices to regulators, clients, or in the context of privilege disputes, a locally processed dictation solution provides a straightforward compliance narrative. The firm can document that its dictation tool processes all audio locally, that no dictation data is transmitted to any external server, and that this can be independently verified through network monitoring.

> **In the event of a privilege challenge, the firm's ability to demonstrate that no privileged dictation data ever left the attorney's device is a complete defence to any argument that the use of dictation technology constituted a waiver. The firm need not argue that its precautions were reasonable; it can demonstrate that the category of risk never existed.**

8. Conclusion

The attorney-client privilege exists to serve a vital function in the administration of justice: ensuring that clients can communicate freely with their lawyers, and that lawyers can develop their analysis and strategy without fear of compelled disclosure. The technology an attorney uses to capture and record legal work must be evaluated against this purpose.

Cloud-based dictation services, whatever their other merits, introduce a structural tension with the privilege framework. They require the transmission of the attorney's spoken words — words that may constitute privileged communications, opinion work product, or both — to servers operated by third parties. That transmission creates questions about waiver, voluntary disclosure, vendor access, metadata generation, and subpoena risk that do not have easy answers and that place the burden of proof on the attorney asserting the privilege.

The standard imposed by the Model Rules is not perfection. It is reasonable precautions, reasonable efforts, reasonable care. But when a technology exists that eliminates the category of risk that those precautions are designed to address — that processes dictation entirely on the attorney's own device, with no network transmission, no vendor access, no metadata generation, and no third-party records to subpoena — the reasonableness analysis shifts.

> **The most secure network connection is the one that does not exist. For the attorney who values the privilege and takes seriously the duty of confidentiality, local processing is not one precaution among many. It is the architecture that makes the question of precautions unnecessary.**

Disclaimer: This white paper is provided for informational purposes only and does not constitute legal advice. The analysis presented reflects the author's interpretation of the cited rules, opinions, and case law principles as of the date of publication. Attorneys should consult their own professional responsibility counsel and applicable state bar ethics opinions when evaluating technology choices for their practice. The ethical obligations discussed herein may vary by jurisdiction, and attorneys are responsible for compliance with the rules applicable in their jurisdiction of practice.

© 2026 Air Gap Voice | Public Release