



Satya's practice focuses on technology transactions, intellectual property, and commercial agreements for startup and growth-stage technology companies across AI, robotics, cloud computing, and emerging tech sectors.

*Email: snarayan@gcalaw.com
LinkedIn Newsletter: [Techline Brief](#)
Connect on LinkedIn*

CONTENT FROM AI NOTETAKING APPS CAN BE A GOLD MINE FOR OPPOSING COUNSEL

[Satya S. Narayan](#)

July 28, 2026 (Updated)

Subscribe to my newsletter: [Techline Brief](#)

Most people think AI notetaking apps create a transcript. In reality, they often create a searchable record of who spoke, what decisions were discussed, what action items were assigned, and in some cases even how participants were perceived to feel during a conversation. In litigation, that can become a very different kind of record than many users expect, a theme I explored from a different angle in an earlier piece on AI chat tools and discoverability [here](#).

Why AI Notetaking Records Will Matter in Litigation

Under *Federal Rule of Civil Procedure 34(a)(1)(A)*¹, electronically stored information, commonly called **ESI**, includes writings, recordings, images, and other data or data compilations, in essentially any digital format or medium. That definition is written broadly on purpose, and it reaches AI notetaking output squarely. Transcripts, including metadata such as speaker labels and timestamps, and analytical outputs such as summaries, would all fall within the ESI definition. As such, AI outputs of these apps are discoverable so long as they are relevant to a claim or defense and proportional to the needs of the case under *Rule 26(b)(1)*². Just like emails and text messages, AI-generated content can become a liability, and potentially more so given the comprehensive nature of transcriptions and AI-generated analysis of internal deliberations, examples of which are discussed below.

Each state also has statutes, protocols, and/or guidelines regarding ESI. The *California Electronic Discovery Act*,³ defines ESI broadly enough to capture the same categories as federal law, and *Rule 3.724(8)* of the *California Rules of Court* requires parties to meet and confer over ESI preservation and scope before the first case management conference.

I am not aware of any court that has yet compelled the discovery of the output from AI notetaking apps; the doctrine, however, is not novel. It has already been applied to structurally identical categories of data on other platforms.

In *Benebone LLC v. Pet Qwerks, Inc.*,⁵ a dispute arose over whether Slack messages should be included in discovery at all. A federal magistrate judge in the Central District of California allowed it, finding that reviewing and producing Slack data was generally comparable to producing email. Benebone argued that production would be unduly burdensome because its Slack workspace contained approximately 30,000 messages and would be costly to review. The court nevertheless concluded that volume alone was not enough to place the data beyond discovery; it credited testimony that targeted searches and specialized review tools could keep the process proportional to the needs of the case.

Notably, it was the defendant, Pet Qwerks, who moved to compel production of the plaintiff Benebone's own internal Slack messages, where its marketing director, COO, and CEO discussed product design and business strategy. That may seem counterintuitive, but it reflects a basic feature of litigation strategy: it is not just the plaintiff who can seek discovery to build their case. Discovery is just as important to the defendant, who can use it to find material in the plaintiff's own records that undercuts the very claims being asserted against the defendant. The case settled within weeks of the ruling compelling production.

A seasoned litigator would likely argue that a speaker-labeled, time stamped transcript from an AI notetaking app is not materially different from other forms of searchable workplace communications, such as Slack and email messages. It preserves candid internal discussions in a reviewable format, and that candid, unfiltered discussion can become exactly the kind of record either side in a dispute may want to use against the other.

Beyond the Transcript: The Discovery Trail Left by AI Notetaking Apps

Most executives adopt these tools for two things: an accurate transcript so nobody has to take notes, and a summary so they don't have to read the transcript. In practice, many meeting platforms generate additional layers of information beyond those two, although the specific features vary by product, subscription level, and enabled integrations. Depending on how the notetaking tool is configured, a single meeting may produce substantially more ESI than just a record of what was said.

Speaker-attributed, searchable transcripts. Modern AI notetaking tools do more than create a transcript. Fireflies.ai states that it automatically identifies different speakers and allows users to search meeting content down to specific sentences and timestamps.⁶ Otter.ai similarly advertises speaker recognition and searchable meeting records.⁷ On both platforms, speaker labels and basic keyword search are available at the free tier; searching specifically by speaker requires a Pro plan or above. Without an AI notetaking app, a litigant most likely would not have gotten a recording at all. Most business meetings are not recorded, so the pre-AI baseline is often no contemporaneous record beyond sparse notes, emails, or memory, and even the recorded meetings that have long existed, earnings calls and, in some cases, board meetings, still meant sitting through raw audio with no speaker labels or search function. An AI notetaking app changes that baseline: the

transcript it produces already comes organized by speaker and timestamp, arriving in a form ready-made for review software rather than requiring days of manual listening.

Action items and assigned responsibilities. AI meeting tools can also convert discussion into structured records of follow-through. Fireflies.ai states that it automatically identifies action items and associates them with relevant speakers.⁶ Otter.ai similarly advertises automatic capture of action items and follow-ups.⁷ In a later dispute, that information may be used to show not only that an issue was discussed, but also who was tasked with addressing it and when the assignment was made.

Analysis of the deliberation itself. Some platforms offer more features at additional subscription levels. Fireflies.ai's site describes a conversation-intelligence layer that compares trends and tracks how frequently a topic, objection, or issue comes up across meetings over time.⁸ Otter.ai's comparison of its product against a competitor states that its AI Chat feature lets anyone on a team ask questions across months of meetings and receive an answer with attribution back to the original transcript.⁷ In organizations that enable these features, the resulting dataset can make it possible to trace how discussions evolved across multiple meetings leading up to a decision.

A record that schedules its own follow-up. Some tools do not stop at noting that a next step was discussed. Motion's help documentation states that once a user approves an action item extracted from a meeting, the task is created and automatically scheduled on the relevant calendar, and rescheduled automatically if details later change.⁹ That creates a direct, dated link between what was said in a meeting and a calendared follow-up task to act on it. The result is a discoverable trail showing not just that a task was assigned, but exactly when it was created and whether it was later rescheduled, reassigned, or quietly dropped.

Tone and sentiment, scored. Perhaps least understood, some meeting-intelligence platforms analyze how participants communicated, not merely what they said. This typically sits at higher level subscription tiers offered by AI notetaking tools. For instance, Fireflies.ai's site states that its platform can analyze talk time, sentiment, and conversation trends, with speaker insights and team-level metrics, features available on its Business plan and above.¹⁰ That is not a transcript of what was said. It is an AI-generated characterization of how someone said it, and how much space they occupied in the conversation.

Consider a termination meeting in which the notetaking app assigns the manager a 'hostile' tone score, or a cross-meeting analysis that repeatedly characterizes a team lead as 'dismissive' whenever engineers raised a particular safety concern. Whether or not those characterizations are accurate, they become part of the record the moment they are generated.

In disputes turning on intent or hostility, admissibility of sentiment scoring may face challenges based on the grounds of reliability and methodology, but it may still be discoverable, and that alone can give a litigant valuable information:

- It can guide depositions. A sentiment score flagging a tense or hostile exchange tells opposing counsel exactly which meeting, which moment, and which participants are worth deposing closely, even if the score itself never comes into evidence.

- It can support settlement leverage. A party may prefer to settle rather than risk a jury hearing about a scored "hostility spike" during discussions of a termination or a contract dispute, even if that score would ultimately be excluded at trial.
- It can serve as a roadmap for further discovery, pointing to related documents, calendar entries, or witnesses that might otherwise not have surfaced.

None of this is generated with litigation in mind. It is marketed as productivity: better search, better follow-through, better visibility into how teams communicate. But every one of these layers may constitute ESI under Rule 34(a)(1)(A),¹ and by linking notes to calendars, documents, or project tools, these applications can create additional discoverable trails that many users and meeting participants may not realize are being created.

Does a Consent Problem Change Discoverability?

Some readers might be asking whether the output of a notetaking app that transcribed a meeting without proper consent under state recording and privacy laws is even discoverable. Whether the notetaking app was used with proper consent from every participant likely does not change whether the resulting transcript, summaries, or metadata are discoverable. Rule 26(b)(1)² turns on relevance and proportionality, not on whether the recording was lawfully made. The idea that improperly obtained evidence is automatically shielded from use is a criminal law concept tied to the Fourth Amendment, and it does not extend to civil discovery between private parties. An opposing party in an antitrust matter or a products liability case may likely still seek production of the transcript, summaries, and metadata if it is relevant and proportional to the case at hand.

The same logic applies to biometric data specifically. Some notetaking apps generate voiceprints, biometric identifiers derived from a speaker's unique vocal characteristics, as part of speaker-recognition features, and litigation is already testing whether that practice violates consent requirements. *Fricker v. Fireflies.AI Corp.*,¹¹ filed in March 2026, alleges that Fireflies' speaker-recognition feature creates voiceprints of meeting participants, including non-users who never consented, in violation of Illinois' Biometric Information Privacy Act, which classifies voiceprints as biometric identifiers subject to strict written-consent requirements;¹² the case remains pending, and at this stage the allegations are unproven. Even if such a consent violation is ultimately established, that may not make the resulting voiceprint immune from discovery. Courts have compelled production of comparably sensitive data before: in *Bartis v. Biomet, Inc.*,¹³ a federal court granted in part a motion to compel a plaintiff's Fitbit data in a products liability case, ordering production of step-count records while permitting redaction of heart rate, sleep, and location data as not relevant, and the court was explicit that objections to the data went to its weight, not its discoverability. The more likely barrier is Rule 26(b)(1)² proportionality: if speaker identity is not contested, counsel has a real argument that producing the voiceprint itself is disproportionate to the needs of the case.

Admissibility v. Discoverability

Discoverability and admissibility are different questions. Discoverability asks whether a party can be compelled to produce material during litigation. Admissibility asks whether that material can later be used as evidence at trial, a separate fight that typically happens later in the case. Federal

wiretapping law, for instance, contains a suppression provision that can bar use of an unlawfully intercepted communication as evidence in certain proceedings, even where the same material was properly produced in discovery.¹⁴ That possibility, however, is unlikely to neutralize the underlying discovery risk. The material is still produced, still reviewed by opposing counsel, and still shapes the litigation long before any admissibility fight is resolved, regardless of whether a court later excludes it. And let us not forget that the consent violation might create a separate unauthorized recording and privacy problem, brought by someone actually on the call who did not provide consent to the recording or to the scope of the app's use of that recording. For more on this latter topic, refer to my prior article [here](#).

Accuracy and Reliability Concerns

AI-generated summaries, key points, and action lists may lack the accuracy and reliability typically associated with human-created notes. One journalist testing Google's Gemini as a meeting notetaker found it had assigned her story topics she never agreed to take on, with Gemini's own summary cautioning her to verify its accuracy.¹⁵ Otter.ai's guidance on its summarization feature acknowledges similar failure modes, including outright hallucination, where a summary states something that was never said, and the dropping of qualifiers, where a tentative comment becomes a firm statement.¹⁶

Producing the underlying transcript can expose discrepancies like these, and a party may argue that an inaccurate summary is unreliable and therefore should not be admitted into evidence as irrelevant or misleading. However, an inaccurate summary can still become relevant if a litigant argues that an employee or executive relied on the summary to make a business decision, since the summary itself can become evidence of what that person believed at the time, and a court may admit it on that basis regardless of whether it accurately reflects what was actually said.

What You Can Do

- Vet the tool the way a company would vet any vendor before deployment. Ask what metadata the app generates and whether that generation can be limited. Ask what analytical features, including sentiment analysis or speaker recognition, the app produces, and whether it integrates with other tools, such as calendar applications. Because these features often vary by subscription tier, upgrading a team's plan can silently expand the discoverable footprint a meeting leaves behind, and counsel should be involved in evaluating which tier represents an acceptable level of risk for the organization, not just which tier offers the most functionality.
- Discuss with counsel whether notetaking apps should be expressly excluded from certain meetings, such as those involving legal counsel, personnel matters, or confidential business information. Ask counsel what custom disclaimers about accuracy can be added to outputs.
- Establish a clear company policy on the use of AI notetaking apps, and educate employees on what these tools capture and retain, not just the transcript, but the metadata trail underneath it.
- Work with IT to configure organizational settings for AI notetaking apps. The specific controls available depend on which architecture the tool uses. Bots that join meetings and

apps that integrate directly with a conferencing platform are typically manageable through enterprise admin consoles. Tools that capture audio locally through a browser extension or desktop app generally are not, and require device-management policies rather than in-app configuration.

- Disable automatic joining of bot-based notetaking apps by default, and restrict or disable third-party notetaking integrations at the platform's admin level.
 - Restrict which notetaking apps are authorized for use, and use device-management policies to control which browser extensions and desktop notetaking apps can be installed.
 - Limit which analytical features, such as sentiment analysis or speaker recognition, are enabled at the organizational level, where the tool's architecture allows for that control.
 - Set data retention and deletion schedules for transcripts, summaries and metadata, where centrally enforceable.
 - Maintain a list of approved tools reviewed for their data storage, training, and access practices, the primary control available for locally installed capture tools that otherwise fall outside IT's direct configuration reach.
- Make clear to employees that joining a meeting with a notetaking app should be a deliberate decision, not an automatic one. Employees should evaluate the nature of each meeting, particularly where discussions must remain confidential or legal privilege needs to be protected.
 - Remember that advance, explicit, app-specific consent may be required under state recording consent laws and other privacy laws. For more on this point, refer to my prior article [here](#).

The safest assumption is that AI notetaking output, including transcripts, metadata, summaries, analytics, and related records, will likely be treated the same way courts already treat comparable forms of ESI when relevant and proportional to a dispute.

What looks like a productivity tool today can just as easily become opposing counsel's most valuable exhibit tomorrow.

Sources

1. Fed. R. Civ. P. 34(a)(1)(A).
2. Fed. R. Civ. P. 26(b)(1).
3. Cal. Civ. Proc. Code §§ 2031.010 et seq. (California Electronic Discovery Act).
4. Cal. R. Ct. 3.724(8).
5. *Benebone LLC v. Pet Qwerks, Inc.*, No. 8:20-cv-00850-AB-AFMx, 2021 WL 831025 (C.D. Cal. Feb. 18, 2021) (Order Re Defendants' Motion to Compel Slack Messages).
6. Fireflies.ai, Conversation Intelligence, last visited July 13, 2026.
7. Otter.ai, Otter vs. Fireflies: Which AI Notetaker Is Better?, last visited July 13, 2026.
8. Fireflies.ai, Features, last visited July 13, 2026.

9. Motion, AI Notetaker (Help Center), last visited July 13, 2026.
10. ComputerTech.co, Fireflies.ai Review 2026: Best Meeting Assistant? (Tested), last visited July 13, 2026.
11. *Fricker v. Fireflies.AI Corp.*, No. 1:26-cv-02675 (N.D. Ill. filed Mar. 10, 2026).
12. 740 Ill. Comp. Stat. 14/1 et seq. (Illinois Biometric Information Privacy Act).
13. *Bartis v. Biomet, Inc.*, No. 4:13-cv-00657-JAR (E.D. Mo. May 24, 2021).
14. 18 U.S.C. § 2515.
15. Erin Chan Ding, AI Notetakers Have Privacy Risks That Do More Harm Than Good, HuffPost, last visited July 13, 2026.
16. Otter.ai, How to Use AI to Summarize Transcripts, last visited July 13, 2026.

Notice: This content is provided for general informational purposes only and does not establish an attorney-client relationship or constitute legal advice. It may not be complete, accurate, or current, and you should seek guidance from a qualified attorney before taking any action. This content may have been organized with the assistance of artificial intelligence. The content may change without notice and could be considered Attorney Advertising in certain jurisdictions.