ENCYPHER

Suggestion to the Advisory Committee on Evidence Rules

Draft Rule 901(c), Revised Rule 707, and the Published Amendments to Rules 104 and 902(1)

TO: Carolyn A. Dubay, Esq., Secretary, Committee on Rules of Practice and Procedure, Administrative Office of the United States Courts, One Columbus Circle, N.E., Room 7-300, Washington, D.C. 20544 (RulesCommittee_Secretary@ao.uscourts.gov)

COPIES TO: Hon. Jesse M. Furman, Chair, Advisory Committee on Evidence Rules; Prof. Daniel J. Capra, Reporter; Prof. Liesa L. Richter, Academic Consultant

FROM: Eddan Katz, Head of Policy and Government Affairs, Encypher Corporation (eddan.katz@encypher.com)

RE: Suggestion concerning draft Rule 901(c) and revised Rule 707, as set out in the agenda book for the meeting of October 15, 2026, and the amendments to Rules 104(a), 104(b) and 902(1) published for public comment on August 14, 2026 (Docket No. USC-RULES-EV-2026-0036)

DATE: October 9, 2026

I. Statement of interest, disclosure, and the limits of this submission

Encypher Corporation builds content provenance infrastructure on the standards published by the Coalition for Content Provenance and Authenticity (C2PA). It operates a claim-signing certificate authority and a time-stamp authority whose credentials appear on the C2PA trust lists, and it takes part in the C2PA's standards work: Encypher co-chairs the task force that wrote the specification's provisions for text, published January 8, 2026. We therefore have a commercial interest in how the Committee treats records of an item's origin and modification, and a specific interest in two of the six findings of the security analysis reproduced in the agenda book at pages 139 through 142, which concern time-stamping and credential revocation. We state that at the outset so that the Committee can read what follows with it in mind, and because we have tried to write a submission that is useful whether or not the Committee credits our interest.

The clearest way to establish that is to state what we do not ask for. **We do not ask the Committee to publish Rule 901(c) sooner than it otherwise would.** We do not ask for any presumption, safe harbor, or evidentiary weight to attach to a provenance record. We do not ask the Committee to name any standard, specification, vendor, trust list, conformance program, or product in the text of any rule, and we agree with what we take to be the Committee's instinct that it should refuse to do so. We do not ask the Committee to treat the absence of a provenance record as probative of anything. And we do not contend that a favorable technical result establishes the truth of what an item depicts, the identity of a human author, or the occurrence of a depicted event.

Four of our twelve recommendations, Recommendations 1, 5, 7 and 12, run directly contrary to our commercial interest, and we have marked each of them. Recommendation 1 is the one we press hardest.

ENCYPHER

One further limit on what follows. **We do not ask this Committee to address the preservation or discovery of records of origin.** Whether such a record is electronically stored information subject to the ordinary duty to preserve, and what consequences follow when a party's own processing removes one, are questions for the Civil Rules and for the courts applying Rule 37(e). Recommendation 1 depends on a distinction between a record that is absent and a record that was taken, and we identify that distinction in Section IV rather than asking this Committee to draw it.

On the October 15 panels. The Committee has constituted two panels for the morning of October 15, and the Reporter has published the questions each will be asked (agenda book, pages 154 through 160 and 177 through 182). Encypher is not among the panelists and makes no request to be added. Section XV instead sets out short written answers to the nine questions put to the deepfake panel and to the three questions put to the Rule 707 panel on which we have something to add, so that the Committee has a provenance-side answer in front of it from a party that has disclosed its interest. Where a point would benefit from technical elaboration, we have said so and stopped, and we will supply that elaboration after the meeting in whatever form the Committee or the Reporter prefers.

On the two dates that bear on this filing. The amendments to Rules 104(a), 104(b) and 902(1) approved for publication on May 7 are open for comment on Docket No. USC-RULES-EV-2026-0036 until February 15, 2027, and requests to be heard at the Committee's hearing of January 11, 2027 are due by December 8, 2026. Encypher will file Recommendation 11(a) in that comment period as well as here, and will request to be heard, because it is the one point in this submission that the Committee can resolve without waiting for draft Rule 901(c). The Reporter's memorandum reports that, as of September 15, no comments addressing the Rule 104 amendments had been received on the rules website, and that three evidence scholars had sent him informal comments (page 190).

Finally, a word about the record we rely on. Every characterization of the Committee's own reasoning below is drawn from the agenda book for the meeting of October 15, 2026, including the Reporter's three memoranda of September 15, 2026, the draft minutes of the meeting of May 7, 2026, and the draft minutes of the June 2026 meeting of the Standing Committee; from the agenda book for the meeting of May 7, 2026 and the Reporter's memoranda of April 1, April 5 and April 11, 2026; and from the November 5, 2025 agenda book and December 1, 2025 report. Unless otherwise stated, page references are to the October 15, 2026 agenda book, and line references are to the line numbers printed in the working drafts reproduced there at pages 150 through 152 (Rule 901(c)), 173 through 176 (Rule 707) and 188 through 190 (Rule 104). Where we disagree with the Reporter we say so and engage the reason he gave, rather than the position.

II. Summary of recommendations

No.	Instrument	Recommendation	Location
1	Rule 901(c)	Add a paragraph providing that the absence of data describing an item's origin or history does not by itself satisfy paragraph (1)	Rule text
2	Rule 901(c)	Adopt the Reporter's factor paragraph with a functional chapeau; separate "missing" from "anomalous" metadata	Committee Note
3	Rule 901(c)	Revise the definition in paragraph (5) to turn on the relationship between the output and the event, and re-attribute it	Rule text and Note
4	Rule 901(c)	Retain the notice provision as adopted on May 7, with no good-cause exception	No change
5	Rule 901(c)	Provide that a conflict among independent signals of origin, including a conflict among verification tools, is a fact for the court to resolve	Committee Note

ENCYPHER

No.	Instrument	Recommendation	Location
6	Rule 901(c)	State what a favorable technical result does and does not establish, with the factors the security analysis in the agenda book shows a court will need	Committee Note
7	Rule 901(c)	Address the interaction with Rules 902(11) through (14)	Committee Note
8	Rule 901(c)	Confirm that paragraph (2) does not alter any burden of persuasion	Committee Note
9	Rule 707	Retain the three new analysis bullets; distinguish records of system operation from records traveling with the output	Committee Note
10	Rule 707	Permit an integrity foundation to be established by certification, as one of the "exceptional circumstances" contemplated by subdivision (b)	Committee Note
11	Rules 104, 902(1), 902(14)	Resolve the standard for a contested Rule 902 item in the Rule 104 Note; correct one Note citation; clarify Rule 902(14)'s reach	Committee Notes
12	Rules 707, 901(c)	Confirm that a certification-based foundation is without prejudice to the Confrontation Clause in a criminal case	Committee Note

Consolidated proposed text is at Annex A. Consolidated proposed Note language, keyed to the line numbers of the October working drafts, is at Annex B. A comparative analysis of statutory definitions, responsive to the question the Reporter posed in April and to panel questions 5 and 6, is at Annex C.

III. On whether an amendment is warranted now

We do not take a position on whether the Committee should publish. We do offer five observations about the record on which that decision will rest, the first of which is unfavorable to the position we take on the merits.

First, and against our interest. The Reporter's April memorandum states that "15 of the responding judges have encountered deepfake issues; and the majority of respondents favored the issuance of a rule on deepfakes" (May 2026 agenda book, page 215), and the Chair summarized the survey for the Committee in May in the same terms (page 20). Both statements are accurate. The Committee should also have in front of it that the majority is drawn almost entirely from respondents without deepfake experience. Among the 889 inexperienced judges who answered Question 9, 487 (55%) favored amendment. Among the 15 judges who had actually encountered a deepfake challenge, 10 (67%) said the rules should not be amended (May 2026 agenda book, Tables 6 and 8, pages 230 and 232). The judges closest to the problem are the least persuaded that a rule is needed. We think the Committee should weigh that inversion, and we think its continued caution is well founded.

Second, and in our view the strongest thing in the survey, unremarked in the summary of it. The Federal Judicial Center's PACER follow-up records that "an informant pled guilty to obstruction of justice after providing false audiovisual evidence to the United States," that "[t]he falsified evidence had resulted in an indictment of an innocent person in an earlier case," and that "[t]here was no court decision about the authenticity of the evidence in the docket of the earlier case, which was dismissed on a motion of the United States with no explanation," the details emerging only at the informant's sentencing (May 2026 agenda book, page 229, citing *United States v. Hernandez*, No. 3:25-cr-2384 (W.D. Tex.), sentencing transcript, ECF No. 28). Magistrate Judge Matthewman's article, reproduced at pages 142 through 145, has now drawn the same case to the Committee's attention from the outside: "if this happened in a federal criminal case, it can happen in any case."

ENCYPHER

That case is invisible to the survey's central question. No litigant challenged an item as a deepfake; the fabrication succeeded, an innocent person was indicted, and the record contains no authenticity ruling at all. If fabricated audiovisual evidence is capable of producing an indictment without ever generating a docketed authenticity dispute, then the number of docketed authenticity disputes is not a measure of the problem. We do not say this proves the rule is ripe. We say the survey measures adjudicated challenges, and *Hernandez* is the case that shows why that is not the same thing.

Third, on the categorical question. One member "remained unconvinced that the issues presented by alleged deepfakes in the courtroom are materially different from the issues presented by more traditional forgeries," noting that "litigants have had to call experts to authenticate signatures after claims of forgery" (page 22). We think that objection is right on its own terms and is not answered by asserting that fabrication has become easier. The difference is not in the forgery. It is in the record.

A forged letter was a physical object with a physical custody trail. A doctored photograph had a negative. Born-digital media routed through ordinary distribution and filing pipelines frequently arrives with no record of its origin at all, because the pipeline removes it. The Committee's own materials make the point three times over: the May agenda book records that "the vast majority of photos submitted as evidence today carry no provenance data at all" and that "[a] photo scraped from social media, or submitted by email with no chain of custody, is just pixels" (May 2026 agenda book, page 208), and the practitioner guidance reproduced in the October book states that "social media platforms like TikTok, Facebook, and Instagram routinely recompress uploads and strip metadata" (page 138). That is a change in the evidentiary substrate rather than in the sophistication of the liar, and it is what makes an authentication rule a different problem in 2026 than a forgery rule was in 1975.

Fourth, the category the survey cannot see. The survey asks whether fabricated items will be admitted. The mirror-image problem is already in front of courts: authentic items attacked as fabrications. *United States v. Belt*, 181 F.4th 931 (8th Cir. 2026), reproduced at pages 134 through 135, is the federal example now in the agenda book. The defendant challenged surveillance video on the ground, as the Reporter summarizes it, that there were "many unexplained gaps in the recording"; the court of appeals held that the gaps went to weight; and Judge Ericksen, dissenting from the denial of rehearing en banc, wrote that "this is a time in which courts should be more stringent in considering foundation" (page 135). The article by Richard Gilman reproduced at pages 145 through 147 reports that courts "are resisting the argument that deepfake possibility creates blanket authentication uncertainty that can excuse liability for recorded statements," the best-known instance being the 2023 suggestion by counsel for Tesla that a recorded statement by its chief executive might be a deepfake, which the court rejected (*Huang v. Tesla, Inc.*, No. 19CV346663 (Cal. Super. Ct. Santa Clara Cnty. Apr. 27, 2023)). Courts elsewhere have met the same move. The Victorian Law Reform Commission's report *Artificial Intelligence in Victoria's Courts and Tribunals* (October 2025) records that New South Wales police unsuccessfully resisted release of body-worn camera footage on the ground that it could be manipulated given the rise "of artificial intelligence and 'deep fake' technology" (*Higgins v Commissioner of Police, NSW Police Force* [2024] NSWCATAD 175, [23]), and that the Ontario Superior Court of Justice admitted surveillance video despite acknowledged gaps in custody, reasoning that "[t]he fact that gaps in the evidence leave open the possibility of a faked or altered video being produced does not necessarily render that possibility a reasonable one absent some evidence tending to make that inference more than a reasonable or far-fetched one" (*R v Gallerno* [2025] ONSC 236, [13]) (Report, paras 5.121-5.122).

We draw two things from this. The Ontario formulation is, in substance, paragraph (1) and Recommendation 2 arrived at by a court without a rule: the showing must be item-specific, and the

ENCYPHER

theoretical possibility of fabrication is not a showing. And the party who bears the cost of an unfounded fabrication claim is the proponent of a genuine item, who in a criminal case is as often the defendant as the government. A rule that calibrates the threshold protects that party. A rule that leaves the threshold uncalibrated does the opposite, and no survey of adjudicated fabrication challenges will record the difference, because in this category the challenge succeeds by being made.

Fifth, on the value of federal guidance whether or not the rule is ripe. State courts are already answering the threshold question, in different ways. Maryland requires that a reasonable juror be able to find authenticity by a preponderance (*State v. Sample*, 468 Md. 560 (2020)); California has declined to require a greater showing for digital images "merely because in theory they can be manipulated" (*People v. Goldsmith*, 59 Cal. 4th 258, 272 (2014)); and the New York Court of Appeals divided four to three over the question this year in *Matter of M.S.* The agenda book adds two state instruments that regulate the problem by placing a verification duty on counsel rather than a standard on the court: Louisiana's 2025 framework and the Connecticut Practice Book provision reproduced at page 135, which the Reporter rightly describes as placing "the burden of doing a deepfake inquiry" on lawyers and pro se filers, "a burden that is not fairly placed on lawyers" who "receive evidence from clients" (pages 135 through 136). State courts, as the Law.com account reproduced at page 149 observes, "often model their own rulemaking after what the federal judiciary is doing." Whatever the Committee does, including nothing, will be read by state rules committees as guidance. We think that is a reason to be careful about what the Note says, which is what Recommendations 1, 2 and 5 are addressed to. It is not a reason to hurry.

IV. Recommendation 1: absence of origin data should not satisfy paragraph (1)

Add to the text of Rule 901(c) a new paragraph, numbered (6) as the draft now stands:

> **(6) Absence of Origin Data.** The absence of data describing the origin or history of an item does not by itself satisfy paragraph (1).

This is the only recommendation in this submission that we urge be placed in the text of the rule rather than the Note. **It is also the recommendation most clearly contrary to our commercial interest, and we ask the Committee to weigh it accordingly.**

The inference the rule will otherwise invite, and the Note now proposes to endorse

Paragraph (1) requires the opponent to "present evidence sufficient to support a finding of fabrication." It does not say what such evidence may consist of, and the Note's guidance has until now been negative: a "broad claim of 'deepfake' is not enough" (Note lines 26-28). Between a bare assertion, which is excluded, and affirmative technical proof of fabrication, which the Reporter and the Department of Justice agree may be unobtainable, there is an obvious intermediate move available to a litigant who lacks the latter: point to what the item does not have.

The Reporter's September memorandum now proposes a Note paragraph listing facts "pertinent to the opponent's burden of going forward," and the first item on that list is "missing or anomalous metadata" (page 153). The fourth is "absence of the native file when a derivative version is offered." We address the list as a whole, and with real appreciation for it, in Recommendation 2. Here we address only the first word of the first factor, because it turns an inference the rule would otherwise merely invite into one the Note would name. If the Note is adopted as drafted, the absence of metadata will be, by the Committee's own enumeration, a fact pertinent to the showing required to warrant an inquiry into fabrication. We think that is wrong, and we think the Committee's own materials show why.

ENCYPHER

That move is attractive precisely because it is easy, and the Federal Judicial Center's data indicates courts are receptive to something in that register. Of the 905 inexperienced judges who answered Question 7, 745 (82%) would require an initial showing. Of the 679 who described what they would require, 283 (42%) said "some type of minimal evidence," the examples given being "an affidavit addressing the dispute" or "technical evidence that the proffered evidence is not authentic" (May 2026 agenda book, pages 231-232). "Technical evidence that the proffered evidence is not authentic" is a category into which an argument from missing metadata fits comfortably, and it will be made.

Once one court accepts absence as the threshold showing, the consequence is structural. Every legacy file, screenshot, platform export, format conversion, accessibility conversion, and re-encoded video becomes presumptively suspect, because none of them carries a record of origin. A rule addressed to fabrication would have created a practical requirement to hold provenance data, which the Committee has not proposed and which, on this record, it would not adopt. The parties least able to satisfy such a requirement are those without institutional workflows: individual litigants, indigent criminal defendants, and pro se parties, who are also the parties most likely to be offering a photograph from a phone rather than a signed capture from a managed system.

The agenda book now states the premise of this recommendation in the Committee's own materials

The practitioner guidance the Reporter reproduces at pages 137 through 139 states the premise of this recommendation in terms. Metadata "can be edited, stripped, or deliberately spoofed with off-the-shelf tools, so internal consistency across multiple fields matters more than any single one." Platforms "routinely recompress uploads and strip metadata." iOS "includes a toggle to remove GPS before sending." The guidance closes its discussion of metadata with the sentence we ask the Note to adopt: "Treat metadata as powerful corroboration where it survives, not as conclusive proof" (page 138). A record that is stripped by ordinary handling is a record whose absence tells the court nothing about the item. What the same guidance treats as informative is *inconsistency*: "Device identifiers should match your client's phone or recording equipment. Timestamps should align with the incident" (page 139). That is the second word of the Reporter's first factor, "anomalous," and we ask that it be kept and that the first word be taken out of the list.

The security analysis reproduced at pages 139 through 142 makes the same point from the opposite direction, about the most durable form of origin record now in use. It finds that "[d]ue in part to expired certificates, some C2PA-signed media have already become unverifiable, sometimes within months." We do not dispute that finding; it describes a real operational problem and Section IX addresses it. But its consequence for paragraph (1) is exactly the consequence of stripping: a genuine item, carrying a genuine record, can present to a court with no verifiable record at all for reasons that have nothing to do with the item. An inference from absence would run against that item.

The recommendation is consistent with the Committee's own premises

It follows from the Reporter's stated rationale. The purpose of paragraph (1) is that a bare allegation must not shift the burden: a broad claim "is not enough to put the court and the proponent to the time and expense of showing that the item has not been manipulated" (Note lines 26-28). An argument from absence is a bare allegation in technical dress. It asserts nothing about the item; it asserts something about the ecosystem the item came from.

It follows from the Committee's own reading of the watermarking literature, and from panel question 3. The Reporter drew from Kassis and Hengartner's *UnMarker: A Universal Attack on Defensive Image Watermarking* the conclusion that "watermarks are not a failsafe" (November 2025 agenda book, page

ENCYPHER

164), and told the Committee that the deepfake problem "was likely to worsen" because embedded authentication data can be removed (December 1, 2025 report, page 307). The third question put to the deepfake panel asks whether "watermarking and labeling programs will effectively limit the democratization of deepfakes" (page 158), and the Reporter's own note on the California statute asks "whether this labelling requirement, and the C2PA technology, can itself be easily evaded" (page 137). We expect the panel to answer, correctly, that marks and manifests can be removed. We agree with that reading, and we think it compels this recommendation rather than cutting against it. If a mark can be removed without a trace, then its absence carries no information about whether it was ever present, and no inference can properly be drawn from absence. The Committee's skepticism about the durability of technical markers is the strongest available argument for making the rule say so.

We would add the distinction that the watermarking literature does not address, and which we set out further in Section XV in answer to panel question 3. Removing a mark and forging one are different attacks with different difficulty. Stripping defeats verification from the mark. It does not enable false attribution, because a cryptographic binding that survives is a binding to particular content signed by a particular key. That asymmetry is why presence and absence should be treated asymmetrically: a surviving record is evidence of bounded facts, and a missing record is evidence of nothing.

Absence will be systematically correlated with the criminal file

There is a further reason absence cannot bear this weight, and it is structural rather than technical. Marking and disclosure mandates are addressed to commercial providers of generative systems, to large distribution platforms and, prospectively, to the makers of capture devices. The California provisions reproduced at pages 136 through 137 are addressed to a "covered provider"; Article 50 of Regulation (EU) 2024/1689, which the Reporter summarizes at page 137, is addressed to providers and deployers of generative systems. Neither is addressed to the systems that generate material in an investigative file, and where a mandate reaches synthetic content generally it has so far carved out law-enforcement use. Article 50(2) requires providers of systems generating synthetic audio, image, video or text content to mark their outputs in a machine-readable format and detectable as artificially generated or manipulated, and then provides that the obligation does not apply "where authorised by law to detect, prevent, investigate or prosecute criminal offences." The deployer disclosure duty in Article 50(4) carries the same exception.

The consequence is that the category of synthetic or AI-processed content most likely to appear in a criminal file is the category least likely to carry a mark or a record of origin. If absence is permitted to satisfy paragraph (1), the inference becomes available against material generated in the ordinary course of an investigation and unavailable against material generated by a commercial system with a signing workflow. That is the opposite of any distribution of the threshold burden the Committee would choose deliberately, and it would be produced by silence rather than by decision. We do not ask the Committee to say any of this in the Note. We state it here because it is the strongest reason we know for putting the negative in the text.

Why the text and not the Note

We recognize the Committee's general preference for guidance in the Note, and the Reporter's observation that "[a]ny detailed list of factors is, of course, better accommodated in the Note rather than in the text" (page 152). We agree about the list. We think this one provision belongs in the text for three reasons.

ENCYPHER

First, the inference is drawn at the threshold, by the party framing a motion, often without the Note in front of anyone. A negative that governs what counts as evidence under paragraph (1) does the work of a limitation on the paragraph and reads naturally alongside it.

Second, the Committee has just shown that it will adopt limiting language of this character. In approving the Rule 902(1) Committee Note on May 7, the Committee kept language, added in response to tribal feedback, emphasizing "that the amendment concerns authenticity only, that other evidentiary requirements continue to apply (such as hearsay rules, relevance, and Rule 403), and that parties may still challenge the authenticity of the covered public documents at trial," and, at the Reporter's suggestion, struck a closing qualifier so that the Note would not "appear to add a condition" (page 32). We ask for the same clarity here, one step higher in the hierarchy, because the inference this forecloses is one the rule itself invites and the Note, as now proposed, would name.

Third, the provision is technology-neutral on its face. "Data describing the origin or history of an item" covers container metadata, a signed manifest, a custody log, and whatever supersedes them. It names nothing that can obsolesce.

The line paragraph (6) does not draw

Paragraph (6) addresses absence. It does not address removal. If the evidence shows that a record of origin existed and that a party or a party's agent removed it, that is not an argument from absence. It is item-specific evidence about the handling of the item, and it bears on paragraph (1) as any other evidence of custody does. Nor does paragraph (6) address the case in which the proponent's own account of the item's creation would have produced a record that the item does not carry: a photograph said to have been taken on a particular device, running a particular application, that writes origin data by default. In that case the absence is an *inconsistency between the item and the manner of creation the proponent asserts*, which is item-specific and which the Note language proposed in Recommendation 2 reaches expressly. The words "by itself" in paragraph (6) are there to preserve both cases. We recommend nothing further, because the consequences of removal belong to the preservation rules rather than to Rule 901.

The second case is no longer hypothetical. As amended, California's AI Transparency Act will require large online platforms, from January 1, 2027, not knowingly to strip provenance data that complies with widely adopted specifications from content uploaded to or distributed on them, to the extent technically feasible, and will require manufacturers of capture devices first produced for sale in the state on or after January 1, 2028 to embed latent disclosures in captured content by default, to the same extent (Cal. Bus. & Prof. Code §§ 22757.3.1(b), 22757.3.3). The codified text differs from the version of the Act reproduced at pages 136 through 137. Once such defaults are in place, a photograph said to have been taken on a covered device, offered without the record the device would have written, presents exactly the inconsistency described above, and paragraph (6) leaves a court free to treat it as one. What paragraph (6) forecloses is the inference against the far larger population of items for which no record was ever written, or from which ordinary handling removed it.

We identify the distinction because the objection to paragraph (6) will be that it rewards stripping. It does not. It denies an inference from a missing record. It does not deny an inference from proof that the record was taken, or from proof that it should have been there.

ENCYPHER

V. Recommendation 2: the Reporter's proposed factor paragraph, with a functional chapeau and four edits

The Reporter now proposes to add to the Note, as its third paragraph, a non-exhaustive list of "facts that would be pertinent to the opponent's burden of going forward" (pages 152-153), and has asked the deepfake panel whether the list "has correctly articulated the relevant factor for triggering the deepfake inquiry" (panel question 7, page 159). We support the addition. It fills the most consequential gap in the current draft: the Note says what does not satisfy paragraph (1), but not what does. The list the Reporter has compiled over four years is, in our view, a good one, and what follows is offered in answer to the question he has put to the panel.

We recommend that the paragraph read as follows. Additions to the Reporter's text at pages 152-153 are underlined and deletions struck through. Our changes are confined to the chapeau, to the first, fourth and seventh factors, and to two additional factors; Annex B sets out the same paragraph clean.

> Under this rule, the opponent's objection must be grounded in facts raising a genuine concern about the offered item. The showing must be specific to the item challenged: it may rest on a gap or break in the item's chain of custody, on an inconsistency between the item and any record of its own origin or modification or between such a record and other evidence, or on a characteristic of the item itself that is inconsistent with the manner of creation the proponent asserts. A general demonstration that fabrication of such items is possible does not satisfy paragraph (1), nor does the absence of a record of origin by itself. The facts pertinent to the inquiry are likely to change as technology develops, but the following is a non-exhaustive list of facts that would be pertinent to the opponent's burden of going forward:
>
> - ~~missing or anomalous metadata~~ metadata that is anomalous, internally inconsistent, or inconsistent with the manner of creation the proponent asserts;
> - inconsistent timestamps or creation dates;
> - unexplained edits or compression artifacts;
> - ~~absence of~~ the proponent's inability to produce the native file when a derivative version is offered, where the proponent's own account indicates that a native file existed;
> - gaps in chain of custody that would raise an opportunity for tampering;
> - a result of a process used to verify a record of the item's origin or modification that indicates alteration of the item or of the record;
> - an anomaly in the path by which the item was acquired;
> - glitches such as unusual audio patterns and problems of synchronization;
> - known use of generative AI tools by the proponent ~~in this case and elsewhere~~ in connection with the item or the case; and
> - independent evidence that is inconsistent with the content of the item.

The chapeau answers the objection the Reporter has stated against his own list

The Reporter's comment on the proposed paragraph states the risk precisely: "by necessarily being underinclusive, the list will lead courts to ignore or downplay other relevant factors," and "factors that raise suspicion right now might never even occur in three years; and of course new factors may arise that aren't on the list" (page 153). That is the same objection he made in November 2025 to the bench-card questions drawn from Judge Grimm and Professor Grossman: "the risk of a heavily detailed note is that it will become outmoded by technological advancements" (November 2025 agenda book, pages 155-156).

ENCYPHER

We think the objection is correct as stated, and that the cure is to tell the court what the factors have in common before telling it what they are. Every item on the Reporter's list, and every item on ours, is an instance of one of three things: a break in custody, an inconsistency between the item and a record or other evidence, or a characteristic of the item inconsistent with the proponent's account of how it was made. Those three categories describe what the court is looking for rather than where to look. "An inconsistency between the item and any record of its own origin or modification" is satisfied today by a file whose container metadata contradicts the claimed device, and will be satisfied in ten years by whatever record then exists, without amendment. A functional description cannot become outmoded by technological advancement because it does not describe a technology. With the chapeau in place, the list can be as concrete as the Reporter wants it to be, because a court that meets a fact not on the list will know whether it belongs.

The Committee also has an existing judicial instrument built on exactly this principle. The *Decision Tree for Evaluating AI-Generated Evidence* released in April 2026 by Hon. Paul W. Grimm (ret.), Professor Maura R. Grossman and Kevin F. Brady under the auspices of The Sedona Conference is organized by the question the court is answering rather than by the tool the party used. That structure is what protects it against the obsolescence the Reporter is right to fear, and it shows that functional description is available to drafters in this area. We describe it here as precedent for the form of the language above, not as authority for its content.

The four edits

First factor. We ask that "missing or anomalous metadata" become "metadata that is anomalous, internally inconsistent, or inconsistent with the manner of creation the proponent asserts." The reasons are those given in Recommendation 1. We add only that the Reporter's own source for the factor says the same: "internal consistency across multiple fields matters more than any single one" (page 138).

Fourth factor. "Absence of the native file when a derivative version is offered" is a different kind of absence from a missing metadata field, because it concerns something within the proponent's knowledge and usually its control, and we think it belongs on the list. We ask only that it be stated as the proponent's *inability to produce* the native file where the proponent's own account indicates one existed, so that it is not read to disfavor the many items for which no native file ever existed in the proponent's hands: a video received by message, a screenshot, a platform export.

Seventh factor. "Known use of generative AI tools by the proponent in this case and elsewhere" is the one factor on the list that is not about the item. Use "in this case" is item-adjacent and we think it belongs. Use "elsewhere" is a propensity inference, and it will be made against the parties most likely to have used such tools for legitimate purposes. We ask that the factor be confined to use "in connection with the item or the case."

Two additional factors. We ask the Committee to add a verification result indicating alteration, and an anomaly in the acquisition path. The first is the case in which a record of origin traveling with the item reports that the item or the record has been changed since it was signed. We have stated it narrowly, as a result that *indicates alteration*, because the security analysis in the agenda book shows that a verification process can also fail for reasons that say nothing about the item, including expired credentials and disagreement among validators (page 141); those failures are addressed in Recommendations 5 and 6 and should not be read as a showing under paragraph (1). The second is the class of fact the Reporter himself offered in May when asked how the showing could ever be made: "the origin of the evidence on a particular device, or its placement among other genuine items" (page 22).

ENCYPHER

The paradigm cases are already in the agenda book

The Reporter's own bracketed note on *Adams v. Anderson*, 2026 WL 1501021 (Ky. Ct. App. May 29, 2026), is the clearest illustration of the standard. The opponent "had requested medical records from the urgent care center and the pharmacy, but both claimed no records existed for the time the child was said to have been taken," and the Reporter observes: "This would surely be a sufficient showing to trigger the deepfake inquiry under Rule 901(c)" (page 134). We agree. It is the last factor on the list, "independent evidence that is inconsistent with the content of the item," and it is item-specific without being technical. Nothing about the photograph's metadata was in issue.

In *Mendones v. Cushman & Wakefield, Inc.*, No. 23CV028772, 2025 WL 2613764 (Cal. Super. Ct. Alameda Cnty. Sept. 9, 2025), which the Reporter reproduced in April and Judge Matthewman discusses at page 142, "the metadata of one of the purportedly-real-but-deepfaked videos showed that the plaintiffs' video was captured on an iPhone 6, which was impossible given that the plaintiff's argument required capabilities only available on an iPhone 15 or newer" (May 2026 agenda book, page 214). That is an inconsistency between the item and its own record of origin, and it is the first factor as we have restated it. The court said of its own capacity in the same opinion that it "does not have the time, funding, or technical expertise to determine the authenticity" of the exhibits (page 142). A court in that position is helped by knowing what kind of thing to look for. It is not helped by a list of tools.

It is also worth noting what made *Mendones* tractable, and *State v. Puloka* with it. In *Mendones* the container metadata survived. In *Puloka* the unenhanced original footage remained available. In each case the court could reach the question because a record or an original had not been removed, which is a matter of circumstance rather than of law. A threshold defined by reference to what the record shows is administrable to the extent that records reach the court. That is a reason for the Note to describe the kinds of facts a court should look for, so that parties know what to preserve and produce. It is also a further reason paragraph (6) is necessary: the same circumstance that makes the question answerable in one case makes absence uninformative in the next.

Two further cases in the Reporter's April memorandum show the cost of leaving the threshold unspecified. In *Burnley v. Valentin*, 2026 WL 767145 (E.D. Va.), a recording was authenticated under Rule 901(b)(1) by a witness familiar with the defendant's voice, a standard the Reporter observed "is too low if deepfakes are involved." In *United States v. Martinez*, 2026 WL 544719 (C.D. Cal.), a claim that recordings were "the product of Artificial Intelligence" was answered with authentication under Rule 901(b)(9), "a less than ideal result" in the Reporter's words. In neither case did the court have a framework for distinguishing an item-specific challenge from an unsupported one. That is the framework the chapeau supplies.

The recommendation answers the Chair's calibration concern

The Chair observed in May that "it might be fairly easy for a defendant to make the requisite showing to shift the burden to the Government. For example, the defendant could produce a sworn statement or affidavit that he was not the person depicted in a Government video" (page 23). The Reporter answered that "the judge would have to find such evidence sufficient under Rule 104(b)." The Federal Judicial Center's data suggests the Chair's concern is well placed: 42% of judges described the threshold as "minimal evidence," instancing "an affidavit addressing the dispute."

A sworn denial that the affiant is the person depicted is evidence about the merits, not evidence about the item. Requiring the showing to be *item-specific* is the calibration that separates a workable threshold from a mechanism for extracting the liar's dividend at the pleading stage, and it is what makes the Reporter's

answer to the ambush hypothetical hold: the government bears only a prima facie burden "unless the defendant first made the requisite showing of deepfakery." A list without a chapeau does not supply that calibration, because an affidavit is not on the list and a court will not know whether its omission was deliberate.

Magistrate Judge Matthewman's proposal, reproduced at page 144, reaches the same place from the bench: the opponent's threshold should be "akin to either reasonable suspicion or probable cause," enough to "prevent cavalier and frivolous boilerplate challenges to legitimate, authenticated, and accurate AI evidence." Reasonable suspicion is, by definition, particularized. The chapeau says what it must be particular to.

VI. Recommendation 3: the definition in paragraph (5)

The definition at lines 16-19 reads:

> "generative artificial intelligence" means a computer system that emulates the structure and characteristics of input data in order to create derived synthetic content, including images, video, audio, text, and other digital content.

We recommend:

> **(5) Definitions.** In this rule, "generative artificial intelligence" means a computer system that creates content, or that alters existing content, so as to depict a person, object, place, or event that the content does not record.

And that the final line of the Note be revised from "The term 'generative artificial intelligence' applied in this rule is derived from Executive Order 14110 (Oct. 30, 2023)" (lines 84-85) to a form that does not rest the rule's vocabulary on a rescinded instrument. Annex C supplies the alternatives.

The Reporter has put the definition to the deepfake panel directly (question 6, page 159), and has asked the Rule 707 panel the adjacent question, "How do you draw the line between AI and other digital information? What about if there is a mix?", using the example of photo-editing software that "is not AI, but which now has AI features" (question 3, page 180). We think the two questions have the same answer for Rule 901(c), and that it is not the answer either panel is likely to give if asked to describe the technology.

Three defects, each identified on the Committee's own record

It is not comprehensible on its face. A member "stated that he was unable to discern what was meant to be covered by the draft definition simply by reading it" (page 17). That is a serious problem in a definition whose only function is to tell a court and the parties what the rule covers.

It may not reach generative alteration of genuine content. The same member "noted that generative AI also can take existing, genuine content and make something else out of it" (page 17). A definition keyed to a system that "emulates the structure and characteristics of input data in order to create derived synthetic content" describes synthesis. Whether it reaches a system that takes an authentic video and alters what a person in it appears to say is, at best, unclear. That is the central case. Ms. Shapiro made the adjacent point, "rais[ing] concerns about the definition of generative AI as applied to criminal matters, where evidence can be created through a combination of generative AI and other AI functions" (page 18).

It is architecture-descriptive, and architectures change. "Emulates the structure and characteristics of input data" is a description of how a class of systems worked when the phrase was written. It is a poor foundation for a rule whose drafters are rightly worried about obsolescence, and it will require

ENCYPHER

amendment for reasons that have nothing to do with evidence law. The Connecticut provision reproduced at page 135 shows the failure mode in its most acute form: it defines generative AI by listing products, "including but not limited to, ChatGPT, Co-Counsel, Google Bard, Neeva, Harvey, Ironclad, Bing, DeepSeek, Grok," two of which, Google Bard and Neeva, had been renamed or shut down before the provision was reproduced in the agenda book. A rule that describes an architecture is on the same road, further back.

Why the recommended formulation is better

It reaches both cases. "Creates content, or that alters existing content" resolves the alteration gap directly, and it does so with a clause the Committee can borrow from enacted federal law: the TAKE IT DOWN Act defines "digital forgery" to include content created "by adapting, modifying, manipulating, or altering an authentic visual depiction" (47 U.S.C. § 223(h)(1)(B)). The Reporter rejected that Act's definition for a reason we think is correct, namely that its "when viewed as a whole by a reasonable person" test is "completely out of place in a rule that requires the *court* to determine that the item is more likely than not a deepfake" (May 2026 agenda book, page 218). That objection reaches the reasonable-person clause. It does not reach the alteration clause, and the alteration clause is the part the draft is missing.

It excludes what the Reporter is right to want excluded. He worried that "a security camera video enhanced using AI-based sharpening filters could be wrongly challenged as synthetic evidence despite being legitimate," and that "AI-powered speech-to-text transcription of court proceedings" should not be swept in (May 2026 agenda book, page 217). Under the formulation above, neither is covered: enhancement that clarifies what the camera recorded does not depict what the content does not record, and a transcript of what was said records it.

But it captures the case where enhancement crosses the line, and the Committee's materials now contain it twice. In *State v. Puloka*, No. 21-1-04851-2 KNT (Wash. Super. Ct. King Cnty. Mar. 29, 2024), discussed again in Professor Bellin's draft article at page 164, AI enhancement of a bystander's cellphone video was excluded because the experts "could not rule out the possibility that the tool 'clarified' the video by fictionalizing it, i.e., by inserting pixels that did not represent what really occurred." In *Linkugel v. City of Los Angeles*, 2026 WL 1430790 (Cal. Ct. App.), reproduced at page 163, photographs of a sinkhole "modified by generative AI" were excluded where the defendant argued the result "was not 'enhancement' but rather resulted in a 'doctored image.'" Detail that no camera recorded is synthetic whatever the tool is called. A definition that turns on the relationship between the output and the event handles both cases correctly; a definition that turns on the tool does not.

It also answers the hybrid-system question the Reporter has put to the Rule 707 panel. *United States v. Andrews*, 172 F.4th 969 (7th Cir. 2026), reproduced at page 163, involved an expert who used photo-editing software to adjust the angle of a fingerprint photograph, and the Reporter observes that to apply Rule 707 "the court would have to determine whether those AI functions were operative when the angle of a photo is adjusted." That inquiry is unavoidable under a definition that asks what the tool is. It is unnecessary under a definition that asks what the output depicts: an angle correction depicts the same print the camera recorded, and is outside Rule 901(c) however it was done. We offer this for Rule 901(c) only. Rule 707 regulates analytical output and needs a definition of a different shape, on which we take no position beyond what is said in Section XV.

It requires no reasonable-person test. The court decides, which is the architecture of paragraph (2).

And it survives architectural change, because it says nothing about how the system works. The survey of more than one hundred statutory definitions of artificial intelligence reproduced at pages 169 through 172

reaches the same conclusion from the other side: "Legislators that predict quantitative metrics or architectures are likely to find the regulation hollow, as development moves in another direction," and the author's recommended alternative is to "regulate something other than the AI system." A definition keyed to the relationship between an output and the event it purports to record is that alternative, applied to the one thing Rule 901 has always been about.

On the Reporter's stated reason for keeping the rule narrow

We support the Reporter, and not the Chair, on the broader question, which is now panel question 5 (page 158). The Chair "queried whether it may make sense to broaden the rule to cover all claims of fakery and not just generative AI, which would have the added benefit of eliminating the need for a definition of 'generative AI,'" and Judge Dever asked "why the burden shifting framework should be applicable only to fakes allegedly made with generative AI as opposed to all alleged forgeries" (page 22). The Reporter's answer, that broadening "would alter the law in many settled areas, including authentication of social media evidence," is in our view decisive. A preponderance gate on every forgery claim would unsettle a large and functioning body of authentication law to no purpose. Recommendation 3 is offered as the narrower alternative: fix the definition so that it covers the harm, rather than widening the rule so that the definition can be dropped.

One premise we think should not be relied on. The Reporter justified confining the rule to generative AI in part on the ground that "fakes prepared by good old AI (as opposed to generative AI) are going to be more easily found out by standard detection tools than fakes prepared by generative AI" (May 2026 agenda book, page 217), and panel question 5 restates the premise as "AI that is not 'generative' is unlikely to create a risk of undetectable deepfakes" (page 158). We do not think the record supports that. The same April memorandum reports that Grok, ChatGPT and Gemini failed to identify non-watermarked Sora videos as AI-generated 95%, 92.5% and 78% of the time respectively, and that where a watermark was present those failure rates fell to 30%, 7.5% and zero (May 2026 agenda book, pages 210-211). In that data, detection performance tracks the presence of a mark, not the generation method. The defensible statement of the premise is that *undetectability by ordinary means* justifies the heightened standard, which is a further reason for a definition that does not name a generation technique.

VII. Recommendation 4: retain the notice provision as adopted

We support the notice provision as it now appears at lines 12-15 of the working draft:

> Unless the court orders otherwise, a party challenging an item of evidence under this subdivision (c) must provide reasonable pretrial notice of the intent to present evidence of fabrication, so that all other parties have a reasonable opportunity to respond to that evidence.

We support the Committee's decisions that notice need not be in writing and that "all other parties" is preferable to "opposing parties" (page 18), and the Note's explanation at lines 71-79 that no specific time period is set because "the deepfake issue is likely to arise in different contexts." **We do not ask for a good-cause exception.** The Reporter set out four options in April and rejected the specific-deadline-plus-good-cause structure on the ground that "a rule with a specific time period is likely to devolve into a rule in which good cause becomes the rule, not the exception" (May 2026 agenda book, pages 218-219). We think that reasoning is sound for a rule of this kind, and we note that Professor Roth's good-cause proposal was directed to Rule 707, where it rests on a different and, in our view, correct rationale: a party "may not know prior to trial that they are going to need to call an expert" (pages 27-28). Two of the articles now in the agenda book also urge early pretrial disclosure: Judge Matthewman calls "[a]n early and fulsome pretrial notice requirement" "fundamental" (page 144), and the Colorado judges list "early

disclosure of AI-generated media evidence and challenges to an adversary's evidence as a deepfake" among the practical steps courts should take (page 145).

We do think the Department of Justice's concern is substantial and that the answer to it depends on Recommendation 2. Ms. Shapiro predicted that notice "would be routinely excused for criminal defendants" (page 18), and that the rule "might also encourage parties to raise deepfake claims, which will delay trials and force the Government to retain experts routinely" (page 23). The Reporter's answer is the right one: the government bears only a prima facie burden "unless the defendant first made the requisite showing of deepfakery," and "[a] bare defense allegation of a deepfake would not suffice" (page 23). That answer is only as strong as paragraph (1) is administrable. A threshold with content is what makes the notice provision unnecessary to enforce in most cases, because most claims will not clear it.

We would also observe, on the asymmetry concern, that the burden-shifting structure operates in favor of the party offering the item, whoever that is. A member "hypothesized a scenario involving a defense deepfake video depicting a cooperator in a damning situation," and the Reporter's response, that "Rule 901(c) would afford the government more protection than it currently enjoys in such a scenario" (page 19), appears to us correct. In the ordinary case the government is the proponent of audiovisual evidence and bears the paragraph (2) burden. In the *Hernandez* fact pattern described in Section III, the government was the recipient of the fabrication and would have been its beneficiary.

VIII. Recommendation 5: a conflict among signals is a fact, not a verdict

Add to the Committee Note:

> Independent signals bearing on the origin of an item may conflict, and different tools applied to the same record may report different results. A conflict among such signals or results is a fact for the court to resolve on the evidence under paragraph (2). Neither the presence nor the absence of any single signal is controlling.

This recommendation is contrary to our commercial interest and we offer it for that reason among others.

The agenda book now contains the finding on which this recommendation rests, in the security analysis at pages 139 through 142. Its third finding is headed "Different Tools Produce Contradictory Results": the authors "tested multiple widely used validators and found that (a) the same image can be labeled valid by one tool and invalid by another, and (b) users may receive conflicting conclusions about the same media, including concerning whether it was AI-generated" (page 141). We do not dispute that finding. We ask the Committee to draw from it the conclusion that a court presented with one verification result has been given one tool's answer, and that the Note should say so.

A second form of the same problem is independent of any defect in any tool. A cryptographically valid record of origin attributing an item to a human editor can coexist, on the same asset, with an embedded mark identifying that asset as the output of a generative model, and can be produced through ordinary editing workflows with no cryptographic compromise. That coexistence is formalized and demonstrated in Alexander Nemecek, Hengzhi He, Guang Cheng and Erman Ayday, *Authenticated Contradictions from Desynchronized Provenance and Watermarking* (Case Western Reserve University and University of California, Los Angeles, arXiv:2603.02378, March 2, 2026), which constructs the workflows and proposes a cross-layer audit protocol that single-layer verification cannot replicate. The construction is concrete and the Committee should have its terms. Five hundred images were produced, each carrying a cryptographically valid manifest asserting only a human editing action and, at the same time, a

watermark identifying it as synthetically generated. The two pipelines took identical watermarked images as input and differed only in the template applied at signing, "reducing the distance between a correctly labeled synthetic image and an authenticated fake to the omission of a single assertion field." Submitted to a deployed public verification tool, the same image was reported as carrying no credential, as AI-generated, or as human-edited with no mention of AI involvement, according only to which manifest was attached; the tool has no mechanism to inspect the watermark layer. The authors identify a further case as more dangerous still: where the embedded mark degrades below the threshold at which it can be detected, the misleading record becomes the only remaining signal. We cite this against ourselves because it is the fact pattern a court will meet, and because a Note that treats one favorable answer as dispositive would be wrong on the published record. A Microsoft Research survey, Jessica Young et al., *Media Integrity and Authentication: Status, Directions, and Futures* (arXiv:2602.18681, February 2026), counts sixty combinations of validation results across manifest, watermark and fingerprint, and five result states; only twenty of the sixty combinations support high confidence that the file is the one that was signed and that its provenance assertions are unchanged.

The practical consequence for a court is that a single answer about origin can be manufactured, and that a court presented with one answer has no way to know whether it was. A court presented with the underlying results, including conflicting ones, does. We would rather the Committee say so now than have the first federal precedent on provenance evidence be a case in which a court was handed a verdict and relied on it.

This is also, we think, the practical form of the concern Judge Erica Yew raised in the material the Reporter reproduced in April (May 2026 agenda book, page 209): "This change in deepfakes and the detection of it and the question about whether it's reliable and whether it should be admissible, has the potential of actually changing that burden of proof without a legal thought process about that." The Reporter's bracketed observation that this "is the question that the New York Court of Appeals was arguing about in *M.S.*" is exactly right, and it is a reason to be explicit about what the court is deciding.

IX. Recommendation 6: what a favorable technical result establishes

Add to the Committee Note:

> Evidence relevant to a determination under paragraph (2) may include a record of the item's origin or modification that is cryptographically bound to the item, together with the result of a process used to verify that record and evidence of custody, source, device, and workflow. In evaluating such a record the court should consider what portion of the item the binding covers; the results of signature and credential validation, including whether the credential had been revoked or had expired and how the verifying process treated that status; the trust policy applied and disclosed; the identity and security of the system that generated the record; the source and verification of any time stamp; the identity, version and configuration of the process used to verify the record; links to source material; preservation; and whether the result can be reproduced by another party. A favorable technical result does not, without more, establish the truth of an assertion, the identity of a human author, or the occurrence of a depicted event.

This is a translation, not a proposal. The premise is already in the Committee's record. The May 2026 agenda book reproduces an account of content credentials as "essentially a cryptographic notary stamp embedded in media at the moment of capture," records that "[t]he National Security Agency endorsed content credential approaches in a January 2025 cybersecurity information sheet," and that "NIST's AI 100-4 report identifies provenance systems like C2PA as among the most promising methods available

ENCYPHER

for tracking the origins of media," concluding: "The concept is exactly right: verify the origin, not the output" (May 2026 agenda book, pages 207-208). The Rule 707 materials reproduce the gap that follows from it: "the extent to which such cryptographically verifiable evidence satisfies admissibility requirements remains to be adequately tested in courtrooms" (May 2026 agenda book, page 123). Magistrate Judge Matthewman's article in the October book asks for the same information in discovery, quoting a three-part inquiry that ends: "Verification: What independent, technical verification exists (metadata, hash values, audit logs)?" (page 144).

We ask the Committee to place beside that premise the limits stated in the same sources and in the October agenda book, because they are the limits a court needs. NIST AI 100-4, *Reducing Risks Posed by Synthetic Content* (November 2024), which the May 2026 agenda book cites for the proposition that provenance systems are among the most promising methods available, also states that a signature "cannot verify that the metadata is accurate; on a technical level, it merely serves a notarization function, indicating that the signer attested to the existence of the metadata at a given time." The security analysis the Reporter now reproduces says the same thing in its own terms: "Provenance describes the history of a file, whereas authenticity concerns, for example, whether the content truthfully represents real-world events. C2PA provides provenance signals, not proof of authenticity" (page 140). We cite these because they are the correct account, and because a Note that adopts the promise without the limit will produce the conflation this recommendation exists to prevent.

The security analysis in the agenda book, and what it means for the Note

Because the analysis at pages 139 through 142 is the most detailed technical material on provenance records the Committee has yet had before it, and because two of its six findings concern functions Encypher performs, we address it directly. We do not dispute that the findings describe the behavior of deployed systems, and we do not ask the Committee to weigh our account against the authors'. We ask the Committee to notice that each finding corresponds to one of the factors in the Note language above, and that the language was drafted so that a court meeting the problem the finding describes would know to ask about it.

The first finding, that time stamps "can be replaced or modified without detection" and that "conforming validators may accept conflicting timestamps as valid," is the reason the Note should direct the court to "the source and verification of any time stamp" rather than to the time stamp itself. A time stamp is a signed assertion by a time-stamp authority; its evidentiary value depends on who issued it and whether the verifying process checked it, and a court should be told both. The second finding, that "many conforming validators fail to check revocation properly" and that "some conforming validators continue to accept known compromised credentials," is the reason the Note should name revocation and expiry expressly and ask "how the verifying process treated that status." The third finding is Recommendation 5. The fourth, that "certain parts of a file (such as metadata)" can be excluded from protection and "altered without invalidating the signature," is the reason the Note should ask "what portion of the item the binding covers," a question this finding shows is not a technicality. The fifth, that some signed media have become unverifiable "sometimes within months," is the reason the Note should name preservation, and is a reason the Committee should not expect a record of origin to be self-proving at trial years after capture without evidence of how its verifiability was preserved. The sixth, that the conformance program "does not define security requirements for conforming validators," is the reason the Note should ask for "the identity, version and configuration of the process used to verify the record" and should not treat a conformance claim as a substitute for that information. All six findings are at page 141.

ENCYPHER

The authors state that C2PA "should not yet be relied upon for high-stakes uses such as ... legal evidence" (page 139). We would put the point differently but not inconsistently: a record of origin is evidence, not a verdict; its weight depends on facts a court can be told; and the Note's task is to tell the court which facts. Every recommendation in this submission is consistent with the authors' caution, and Recommendations 1, 5 and 6 are expressions of it.

Two limits are worth stating concretely, because they are the ones a court will actually meet. A record of origin says nothing about the use to which a genuine item is put: it would not resolve a dispute about authentic surveillance footage presented in a misleading context, because the record and the item are both exactly what they purport to be. And no record survives re-capture: a photograph of a screen displaying an item is a new item, with its own origin and no relationship to the record that accompanied the original. Neither limit is a reason to exclude such evidence. Both are reasons for the Note to describe what the court is being given rather than what it may conclude.

Two drafting points. First, the final sentence is not a concession we have been pressed into. It is the boundary the Committee drew for Rule 902(1) on May 7 and the boundary the Rule 707 Note draws against Rule 902(13), citing the Advisory Committee Note to that rule for the proposition that a certification "does not preclude an objection that the information produced is unreliable" (Rule 707 Note lines 43-48). We would not want it softened.

Second, we ask the Committee not to use the word "validated," and not to permit a foundation to consist of a badge, a screenshot, or a single conclusory label. The C2PA specification distinguishes manifests that are well-formed, valid and trusted, and separately asks whether the asset itself is valid; each state has a different evidentiary meaning, and the security analysis shows that different tools collapse them differently. Collapsing them into one word invites a court to read a technical result as a verdict on the item, so the paragraph above describes results rather than conclusions. We will supply a short technical annex on the status results after the meeting if the Reporter would find it useful.

X. Recommendation 7: the interaction with Rules 902(11) through (14)

Paragraph (3) applies the rule "to items offered under either Rule 901 or 902." The Note's treatment of that consequential provision is two sentences, at lines 64-66, which identify the risk as extending to "many of the items designated in Rule 902 as self-authenticating," most obviously "newspapers and publications." The tab contains no discussion of Rule 902(13) or Rule 902(14).

We recommend adding to the Committee Note:

> A certification under Rule 902(11), (12), (13), or (14) does not preclude an inquiry under this subdivision. In determining a question under paragraph (2), the court may consider a certified record of the results of a process used to verify data describing the origin or history of the item, together with the other evidence bearing on authenticity.

The two halves point in opposite directions and we ask for both.

The first half is contrary to our commercial interest and is already the Committee's position elsewhere. The Rule 707 Note provides that "it is not enough that [the output] is self-authenticated under Rule 902(13)," because "[t]hat rule covers authenticity, but does not assure reliability" (Rule 707 Note lines 43-45). One of the three post-comment submissions the Reporter notes at page 182 makes the same point from a different starting place, urging that the Committee "help prevent authentication under Rules 901(b)(9) or 902(13) from being mistaken for proof of reliability." The same limitation should be explicit here, so that a certificate is not read as ending a paragraph (2) inquiry.

ENCYPHER

The second half is the question the rule will actually generate. *United States v. Gogic*, 2025 WL 3042350 (E.D.N.Y.), reproduced in the April Rule 707 materials, shows the shape of it. The defendant contended that "artificial intelligence was likely used to create the Chat Spreadsheets, and that metadata reveals the Sky Evidence was modified multiple times before it was produced." The court held that "[s]tanding alone, the Certificate of Authenticity merely supports the contention that the Sky Evidence reflects an extraction of certain information maintained electronically in a European 'electronic storage medium'; it does nothing to establish the origins or reliability of that information" (May 2026 agenda book, page 117). That is a case in which the certificate did not answer the question and the record of origin was the contested ground, litigated without guidance from any rule.

We note, without asking the Committee to resolve it here, that Rule 902(13) and Rule 902(14) are the provisions most likely to carry provenance evidence in practice, that the 2017 Advisory Committee Note to Rule 902(14) already contemplates "other reliable means of identification provided by future technology," and that the point has not been taken up in nine years. Recommendation 11 below addresses it in the narrowest available form.

XI. Recommendation 8: the nature of the paragraph (2) determination

We recommend the Note confirm that paragraph (2) does not alter any burden of persuasion. Suggested language:

> A determination under paragraph (2) is a preliminary question for the court. It does not alter the allocation of the burden of persuasion on any issue in the case, and it does not preclude a party from presenting evidence bearing on authenticity to the trier of fact after the item is admitted.

This restates what the Committee said in discussion. The Reporter confirmed that a party may still present authenticity evidence to the jury after admission, and the Chair analogized to allegedly coerced confessions, "where the court plays a gatekeeping role but a defendant may argue coercion before the jury" (page 23). We think it belongs in the Note because *Matter of M.S.*, 2026 N.Y. Slip Op. 00825, 2026 WL 436359 (N.Y. Feb. 17, 2026), and Judge Yew's question both indicate that courts are already conflating the gatekeeping determination with the ultimate allocation of proof. Judge Troutman's dissent in *M.S.* states the confusion in its sharpest form: "a party opposing the introduction of video evidence need only posit, with no evidentiary support whatsoever, that the video might be a deepfake. The party introducing the video then must disprove that unfounded accusation. How is the party seeking admission to do so, if circumstantial evidence is not a permissible method of authentication?"

We think that question deserves an answer in the Note, and that the answer is Recommendations 1, 2 and 6 taken together. The threshold is item-specific; a record of origin is admissible on the shifted question; and its absence proves nothing.

One observation for the Committee's awareness. The version of draft Rule 901(c) quoted by Judge Singas in *M.S.* is an earlier version with different subdivision headings and without paragraphs (3), (4) and (5). Any discussion of what the New York Court of Appeals had before it should use her text rather than the current draft.

XII. Rule 707: Recommendations 9 and 10

Encypher supports the working draft of Rule 707 reproduced at pages 173 through 176, which the Committee unanimously agreed in May to carry forward "with a focus on Version 2" (page 28). We confine ourselves to the two points on which we can add something, and to the Reporter's fifth question to the

ENCYPHER

Rule 707 panel, "Can AI be validated any way other than through an expert?" (page 181), which Recommendation 10 answers in part.

Recommendation 9: retain the three new analysis bullets, with one distinction

The three bullets added in the redrafts, and absent from the published version, now appear at Note lines 69-76:

> Reviewing information about what records the system keeps, what is deleted, and whether the output can be reproduced or audited.
>
> Considering whether the process or system has been subject to evaluation by entities independent of the developer, and whether the opponent and independent evaluators have access to the system, e.g., through meaningfully available research licenses.
>
> Considering whether the opponent has received sufficient information about how and why the process works, including information analogous to that which would be disclosed if the conclusion were that of a human expert.

We support all three without qualification, including the second, which is directed at companies such as ours. The American Civil Liberties Union's comment warns that "documentation often comes from the company that sells the device or software to law enforcement, or from some other entity with incentive to preserve their relationships with the government, and thus may not represent the sober reflections of an independent and therefore qualified expert." We think that is a fair description of a real risk, and the article by Gregory Schwartz reproduced at pages 167 through 168 adds the historical record behind it: "Each time, courts held that source code disclosure was unnecessary to sufficiently vet the software, and each time, serious errors were found when the code was eventually released." We also endorse the position the Reporter quoted from Judge Grimm and Professor Grossman in April: "if a party intends to use AI-generated evidence, it cannot be allowed to do so while simultaneously objecting to any discovery by the opposing party and thereby preventing them from evaluating and challenging the evidence" (May 2026 agenda book, page 163).

One distinction we recommend the first bullet draw. Records that a system keeps *about its own operation* and records that travel *with the output* raise different problems. The former are in the developer's or proponent's exclusive control and are obtainable only through discovery, with all the proprietary-access difficulties the comments describe. The latter accompany the file and are available to any party holding it. Suggested revision:

> Reviewing information about what records the system keeps, what is deleted, and whether the output can be reproduced or audited, distinguishing records retained by the system or its operator from records that accompany the output and are available to any party in possession of it.

This matters for a reason several commenters raised and none resolved. Equality of arms was pressed by the ACLU, by the National Association of Criminal Defense Lawyers, and by Ms. Owens on behalf of the Federal Defenders. A disclosure regime that depends entirely on discovery against a developer favors the party with resources. A record that travels with the output does not. Making the distinction in the Note tells courts which kind of inquiry they are conducting. *United States v. Belt*, discussed in Section III, is the authentication-side illustration: the dissent's complaint, as the Reporter summarizes it, that the prosecution "made no showing about the surveillance taping system" (page 134) is a complaint about the first kind of record, and no record of the second kind was available to supply it.

ENCYPHER

Recommendation 10: permit an integrity foundation by certification

Comment 0012 in the Rule 707 docket proposed Note language which we endorse and ask the Committee to adopt in substance:

> In applying Rule 702(a)-(d) to machine-generated evidence under Rule 707, courts may consider threshold integrity factors necessary for meaningful adversarial testing. Such factors may include whether the proponent can provide tamper-evident records sufficient to detect missing or altered inputs; identify the system, model version, and execution context that generated the output; and permit independent verification of completeness and provenance.

We note the Reporter's own assessment that this commenter "correctly recognizes that authenticity mechanisms, including Rule 902(13), do not establish reliability."

We ask the Committee to add the one element the comment does not address: that where such records exist and accompany the output, the integrity foundation they supply may be established by certification rather than by live testimony. This is the only commercially material request in this submission, and we make it separately so the Committee can weigh it separately. It concerns only evidence the proponent acknowledges is machine-generated, so it raises no question about deepfakes, no question about ripeness, and no question about burden-shifting. It sits inside the architecture of Rule 902(13) as that rule already exists. It is also subject to Recommendation 12, which we ask the Committee to adopt together with it and not separately.

The request fits the structure the Committee has chosen. Subdivision (b) of the working draft provides that "[o]rdinarily, the proponent must provide an expert," but that "in exceptional circumstances, the court may consider other proof to find that the output satisfies the requirements of (a)(2)-(4)," and the Note offers "convincing validation tests and objective studies indicating a low rate of error" as examples (Rule 707 lines 10-13; Note lines 58-60). Professor Roth observed in May that the word "provide" was presumably chosen so as not to require live testimony in every case, since "litigants may offer an expert's conclusions through an affidavit in current Daubert hearings" (page 28). An integrity foundation by certification is a narrower instance of the same idea. It does not validate the model; it establishes the bounded, reproducible facts that any validation must assume: which system and version produced the output, whether the inputs and the output have been altered since, and whether the result can be re-run. Those facts are the precondition for "meaningful adversarial testing" in comment 0012's phrase, and they are the kind of fact for which a certificate is an appropriate vehicle because any party holding the output can check it.

On the Department of Justice's objections to the (b) escape hatch, two observations. The Department's letter of February 13, 2026, argues that "a 'general' or 'independent' validation (by an academic institution, for example) to determine whether a forensic instrument or software 'works' is meaningless because it does not test the application's fitness in a specific, operational testing environment with its own unique conditions, parameters, and intended use or application," and that "for machine learning, outputs can never be perfectly replicated because the machine is continuously learning." We think both propositions are correct about model behavior and inapplicable to a record of origin, which asserts bounded historical facts about a particular file rather than the output of a continuously learning process. A hash either matches or it does not, and the match is reproducible by anyone with the file.

We also agree with the Department, expressly, on a point in its favor and ours. Its letter argues that metadata "often is so accessible that lay witnesses can look at it and explain its meaning to a case. Thus, it should not be subject to Rule 707" (May 2026 agenda book, page 182). We think that is right, and we

think it follows that records of origin and modification are outside Rule 707 and squarely inside Rules 901 and 902, which is where this submission has placed them.

XIII. Recommendation 11: the published Rule 104 and Rule 902(1) amendments, and Rule 902(14)

The amendments to Rule 104(a) and (b) were approved for publication on May 7 without opposition, the amendment to Rule 902(1) with one member opposed, and the Standing Committee approved both in June without opposition (pages 32, 41 and 76-77). They are open for comment on Docket No. USC-RULES-EV-2026-0036 until February 15, 2027, which is before draft Rule 901(c) will be published. Two points arise from reading the Rule 104 and Rule 901(c) tabs together, and one further point concerns Rule 902(14). We raise them here because they bear on Rule 901(c). We will also file paragraph (a) below in the Rule 104 comment period, because it is the one point in this submission the Committee can act on within a docket that is already open, and we do not want it to depend on the Committee reaching draft Rule 901(c).

(a) Rule 901(a) is preserved; Rule 902 may not be. Footnote 4 of the Reporter's April 1 memorandum states that "[a]ll other rules applying the 'sufficient to support a finding' standard are unaffected by a change to Rule 104(b), because each sets forth the standard without referring back to Rule 104(b). See Rules 602 and 901(a)" (May 2026 agenda book, page 355). The September memorandum goes further in response to Professor Imwinkelried and offers, if "in an excess of caution" the Committee thinks it appropriate, a sentence for the Note: "Nothing in this amendment is intended to affect the admissibility requirements under Rules 602 or 901. These are now examples of rules 'otherwise provided'. See Rule 104(a)" (pages 195-196). We agree with both statements, we support adding the sentence, and we do not contend that authentication under Rule 901(a) migrates to the amended Rule 104(a).

Rule 902 is different, and the proposed sentence makes the difference sharper by naming Rule 901 and not Rule 902. Rule 902 provides that listed items "require no extrinsic evidence of authenticity in order to be admitted." It states no standard of proof. It therefore does not "otherwise provide[]" within the meaning of the amended Rule 104(a), and it will not be among the "examples of rules 'otherwise provided'" under the Note as the Reporter proposes to revise it. Yet the authenticity of a Rule 902 item may always be contested. The published Note to the Rule 902(1) amendment says so ("parties may still challenge the authenticity of the covered public documents at trial," lines 24-26, page 185), and describes what self-authentication confers as "the same presumption of authenticity" the other listed documents enjoy (lines 12-14, page 184), without saying what showing rebuts that presumption. The question the amendment leaves open is what standard governs the resolution of that contest. On the face of amended Rule 104(a) the answer would be a preponderance, which would mean that the Rule 104(a) amendment does for every contested self-authenticated item what draft Rule 901(c)(2) does only for items challenged as generative-AI fabrications, and that paragraph (3) of the draft, which expressly reaches Rule 902 items, would overlap with it on the same facts under a different trigger.

We do not suggest which way the Committee should resolve this. We ask only that the Rule 104 Committee Note resolve it, in either direction, and the sentence the Reporter has now proposed is the natural place: it need only say whether Rule 902 is, or is not, among the rules that "otherwise provide." The Reporter's memoranda do not discuss Rule 902 or self-authentication, so this appears to be a gap rather than a disagreement. We would put the priority more strongly than the drafting point suggests. Of everything in this submission, this is the item with the widest reach and the shortest timetable. Resolving it settles the standard of proof for every contested self-authenticated item, whatever the Committee

eventually decides about deepfakes, and it can be resolved in a Note to an amendment that is already published.

(b) One citation in the draft Rule 104 Note will become inaccurate. Lines 21-24 read: "Some situations call for a lower standard of proof for a preliminary question ... See Rules 104(b), 602, 901 and 1008." If Rule 901(c) is ever adopted, Rule 901 will contain both a lower standard, in 901(a), and a higher one, in 901(c)(2). We recommend "Rule 901(a)," here and in the sentence proposed at page 196.

(c) Rule 902(14) and "future technology." The 2017 Advisory Committee Note to Rule 902(14) states that the rule is flexible enough to allow certifications through processes other than hash comparison, including other reliable means of identification provided by future technology. That invitation has not been taken up, and the deepfake tab does not mention the rule. We recommend, whenever the Committee next has occasion to address Rule 902(14):

> A process of digital identification within the meaning of this rule includes verification of a record of an item's origin or modification that is cryptographically bound to the item. A certification under this rule establishes only the identity of the copy and the technical results described. It does not establish the truth or completeness of any assertion, the identity of a human author, the reliability of machine-produced information, or admissibility under any other rule.

United States v. Dunnican, 961 F.3d 859 (6th Cir. 2020), illustrates the division of labor the second sentence preserves: a qualified person's certificate and a digital fingerprint authenticated a forensic extraction under Rule 902(14), while separate evidence linked the device and its contents to the defendant. Vermont has gone further by statute, making a blockchain record self-authenticating under its Rule 902 on the declaration of a qualified person, with presumptions we do not ask this Committee to adopt (12 V.S.A. § 1913). We also note that the April Rule 707 materials record Professor Imwinkelried's view, as the Reporter put it, that "[t]here is a strong case that 901(b)(9) does not belong in Rule 901 and that the validation of a process or system should not be governed by Rules 104(b) and 901(a)" (May 2026 agenda book, page 159), and that comment 0056 would amend Rule 901(b)(9) rather than adopt Rule 707. The Committee has an open question about where the validation of a process belongs. A Note clarification on Rule 902(14) is the narrowest answer available to it.

XIV. Recommendation 12: the Confrontation Clause limit on any certification route

We recommend that the Committee Note to Rule 707, and any Note treating certification as a route to an integrity foundation, state that nothing in the rule affects a criminal defendant's rights under the Confrontation Clause. Suggested language:

> Nothing in this rule affects the application of the Confrontation Clause. In a criminal case, the sufficiency of a certification to establish a foundation remains subject to the defendant's right to be confronted with the witnesses against him.

This recommendation is contrary to our commercial interest. Recommendation 10 is the only commercially material request in this submission, and this recommendation narrows it in the setting where it would matter most.

The problem the Committee's materials do not reach

A certification-based foundation in a criminal case is a testimonial-statement problem. *Melendez-Diaz v. Massachusetts*, 557 U.S. 305 (2009), holds that a forensic certificate prepared for use at trial is testimonial. *Bullcoming v. New Mexico*, 564 U.S. 647 (2011), holds that the analyst who made the certification, and not a surrogate, is the witness the accused is entitled to confront. *Smith v. Arizona*, 602

U.S. 779 (2024), holds that when an expert conveys an absent analyst's statements in support of his opinion, and the statements provide that support only if true, the statements come into evidence for their truth; it left open whether the statements there were testimonial. The Committee is itself considering an amendment to Rule 703 to flag *Smith* (pages 33-34 and 84); the same flag belongs on any certification route under Rule 707.

A certificate attesting to the results of a verification process run on an item offered against a defendant falls within that line if it was prepared for litigation. It is arguably outside that line where the underlying record was generated in the ordinary course at the moment of capture, the certificate reports only the result of a process, and any party in possession of the item can re-run that process and obtain the same result. That distinction has not been tested in any court, and we do not ask the Committee to resolve it.

The October agenda book now contains the scholarly argument for caution on exactly this point. Benjamin Welton's *Meaningful Machine Confrontation*, 76 Stan. L. Rev. 845 (2025), reproduced at pages 168 through 169, argues that a stronger reliability foundation must be established for machine-learning evidence offered against criminal defendants. We agree, and nothing in Recommendation 10 is meant to lower that foundation.

What we do ask

We ask only that the Note not be taken to have resolved it by implication. A vendor requesting a certification route has an obvious interest in the Note being silent on this point; we ask that it not be. The American Civil Liberties Union, the National Association of Criminal Defense Lawyers, the Federal Defenders and the Innocence Project have all filed in the Rule 707 docket on disclosure and equality of arms, and this is the point on which a certification route is most exposed to their objection. We think the objection is sound. We would rather the rule say so than have Recommendation 10 adopted in a form that must be litigated out of the criminal courts one case at a time.

We state the corollary in our favor as a limit and not as an argument. Where a record of origin was created at capture, is reproducible by any party in possession of the item, and is offered only to establish the identity of a copy and the technical results described, the Confrontation objection has considerably less force, because the defendant is not being asked to accept an absent analyst's conclusion. That is the case for which Recommendation 10 is drafted, and the Note should say so rather than leaving the criminal application to be inferred from silence.

We record one gap in our own preparation, because the Committee should discount our confidence accordingly. Our analysis of this line rests on secondary treatment of *People v. Wakefield*, 38 N.Y.3d 367 (2022), and the source-code cases, and we have not completed primary-source work on how *Smith* bears on machine-produced certifications. We would not press Recommendation 10 in a criminal case beyond the boundary stated above until that work is done.

XV. Answers to the questions put to the October 15 panels

The Reporter has published nine questions for the deepfake panel (pages 158-160) and nine for the Rule 707 panel (pages 180-182). We answer the nine deepfake questions in order, and the three Rule 707 questions on which we have something to add. Our answers are those of a party with a disclosed interest, offered so that the Committee has them in writing beside whatever the panels say.

Deepfake question 1. What is the likelihood that deepfakes will flood the courts? We do not know, and the docketed-challenge count cannot tell the Committee: *Hernandez* shows that a successful fabrication generates no docketed challenge, and the mirror-image cases in Section III show that an unfounded

challenge to a genuine item succeeds by being made. The Law.com report reproduced at pages 147-149 quotes a Virginia circuit court's chief judge on what state judges are telling one another: "Nobody has a solution, but everybody's talking about and telling more stories." The question the Committee can answer now is what a court should be told to look for when a case arrives.

Deepfake question 2. How hard is it to determine whether an item is deepfaked, and will that get harder? Detecting a fabrication from the item alone is hard and will get harder; the detection data in the April memorandum shows performance tracking the presence of a mark rather than the generation method (May 2026 agenda book, pages 210-211). The word "detection" runs two questions together. Determining from the pixels whether an item was generated is an inference problem that is losing ground. Determining whether an item is the same item that a particular device or system produced at a particular time is a verification problem, and it does not get harder as generators improve, because it depends on a record rather than on the item's appearance. The rule should not depend on the first and should not foreclose the second.

Deepfake question 3. Are watermarking methods, such as C2PA, sufficient to assure that ordinary people will not be generating undetectable deepfakes? No, and the question contains a conflation the Committee should not carry into the Note. A watermark is a signal embedded in the content itself, designed to survive ordinary handling and to be detected by a tool that knows what to look for. A provenance manifest is a signed record attached to the file, designed to be verified rather than detected, and removed by any process that discards the file's metadata. The security analysis the Reporter reproduces draws the same line: C2PA provides "provenance signals, not proof of authenticity," and "was not designed to determine authenticity" (page 140). Neither technology prevents anyone from generating anything; both can be removed; and neither was designed to assure that unmarked content is authentic. What a surviving, verifiable record can do is establish bounded facts about a particular item's origin and handling, and what a missing record can do is nothing. That asymmetry is Recommendation 1. The California statute, which the Reporter suggests "will make it significantly more difficult to get a deepfake into court" (page 137), will make the *presence* of a disclosure more common and more meaningful. It will not make the absence of one informative, because the unmarked population will continue to include everything authentic, everything produced by systems outside the statute's reach, and everything that passed through a platform that strips metadata.

Deepfake question 4. Are there special problems in regulating deepfake text? Yes, and they cut in favor of the functional approach in Recommendation 2. Detection from the item alone is weakest for text, and the Ninth Circuit's *Al-Qarqani* matter (page 134) turned on a purported newspaper article the court found "fabricated for purposes of this litigation," not on any technical marker. The factors that will work for text are custody, inconsistency with independent evidence, and the proponent's account of creation, which are the three categories the chapeau names. A list of audiovisual markers will not help a court with a document.

Deepfake question 5. Should the rule cover generative AI or all AI? Generative AI, for the Reporter's reason that broadening would unsettle settled authentication law (page 22), with the definition repaired as Recommendation 3 proposes. We would not rely on the premise that non-generative fakes are more readily detectable, for the reason given in Section VI.

Deepfake question 6. Does Rule 901(c) provide an accurate definition of generative AI? No, for the three reasons in Section VI, each stated on the Committee's own record. We have checked the Reporter's April statement that "[t]here is no definition of generative artificial intelligence that I could find in a federal statute" (May 2026 agenda book, page 220), and we believe it is correct: the only federal statutory

definition in this area is of "artificial intelligence," at 15 U.S.C. § 9401(3), the source of Rule 707(e) and of the New York court rule reproduced at page 173. Annex C sets out the enacted state and comparative definitions and explains why we recommend a definition keyed to what the output depicts rather than to how the system works.

Deepfake question 7. Is the list of factors useful and comprehensive? Useful, yes, and we support adding it. Comprehensive, no, and the Reporter is right that no list can be. Recommendation 2 proposes the cure: a chapeau that states what the factors have in common, so that the list can be concrete without being exhaustive, together with four edits, the most important of which is to separate "missing" from "anomalous" metadata.

Deepfake question 8. Would a lawyer certification requirement work instead? No, for the reasons the Reporter gives about the Connecticut provision at pages 135-136. We add one from our side of the problem: a certification requirement asks the lawyer to vouch for an inquiry the lawyer usually cannot perform, while an evidence rule asks the court to decide a question on evidence the parties can produce. Where states adopt certification regimes, as Louisiana and Connecticut have, a federal Note that describes an item-specific showing will be the best guidance those lawyers have for discharging the duty.

Deepfake question 9. How likely is it that expert testimony will be required, and at what expense? Less likely, and less expensive, to the extent the Note tells the court what to look for. Most of the factors on the Reporter's list, and all three categories in our chapeau, are facts a lay witness can establish: where the file came from, whether a native file exists, what the device was, what independent evidence shows. The threshold showing in *Adams v. Anderson* rested on clinic and pharmacy records. The cases in which an expert is required are the cases in which the item must be analyzed from its pixels alone, and the point of a Note that directs attention to records, custody and corroboration is to reduce the number of cases that reach that stage. The equality-of-arms concern raised in the Rule 707 comments applies here with equal force, and a record which travels with the item and can be verified by any party holding it is the one form of technical evidence that does not favor the party with the larger budget. Our offer in Section XVI of a free verifier that returns full results is addressed to this question.

Rule 707 question 3. How do you draw the line between AI and other digital information, and what about a mix? For Rule 901(c), by asking what the output depicts rather than what the tool is; Section VI explains why that dissolves the *Andrews* problem. For Rule 707 we take no position on the definition.

Rule 707 question 5. Can AI be validated any way other than through an expert? Not the model, in our view, and we do not ask the Committee to say otherwise. But the Committee should distinguish validation of the model from the integrity foundation on which any validation rests: which system and version produced the output, whether the inputs and the output have been altered since, and whether the result can be re-run. Those facts are bounded, reproducible and checkable by any party holding the output, and Recommendation 10 asks that they be provable by certification, subject to Recommendation 12. That is the narrowest sensible content for the "other proof" that subdivision (b) contemplates, and it is content the Committee can define now without predicting what "self-validating AI" might come to mean.

Rule 707 question 6. How important is access to the source code? For validation of the model, we defer to the panel and to the comments. For the integrity foundation, the relevant disclosure is the record itself and the means of re-running the verification, which is what Recommendation 9's distinction is designed to surface. An opponent who holds the output and the record can test that foundation without any access to the developer, which is the one part of a Rule 707 inquiry in which equality of arms can be achieved by design rather than by discovery.

ENCYPHER

XVI. Offer of assistance

We make six offers, each of which the Committee can accept without accepting any position we have taken.

A compiled and verified survey of enacted definitions. Annex C is a first cut. We will supply a complete compilation of enacted federal, state and comparative definitions of generative artificial intelligence, synthetic content and digital forgery, including the Connecticut and New York court rules now in the agenda book, with verified statutory text and a drafting analysis, in whatever form is useful, at no cost and with no attribution required.

Empirical work on whether provenance data survives court systems. No published study establishes whether federal or state electronic filing and evidence-management systems preserve or remove records of origin on intake. The question bears directly on paragraph (1), because a court cannot draw an inference from what its own filing system removed, and on Recommendation 1. We are undertaking that testing against named systems and will publish the results whatever they show, including results unfavorable to the technology we sell. We will supply them to the Committee and to the Federal Judicial Center when complete.

A free verifier with full results and no conclusion. We will make available, at no charge and without registration, a verification tool that returns the complete set of status results rather than a summary label or badge, together with a published statement of its limitations and of the trust policy it applies, so that a party without an expert or a budget can read the underlying results in the same form as any other party. We mention it because the equality-of-arms concern raised in the Rule 707 comments is a real objection to the whole category, because the security analysis in the agenda book shows that validators disagree, and because we would rather answer both with a commitment than an argument.

A short technical annex on verification status results, after the meeting. Recommendation 6 asks the Committee not to use the word "validated." If the Reporter would find it useful, we will supply within thirty days of the meeting a short annex describing the distinct status results a verification process returns, what each does and does not establish, and how the six findings of the security analysis at pages 139-142 map onto them, written for a reader who is not a technologist and reviewed against the panel discussion.

A provenance branch for an instrument the judiciary already has. The *Decision Tree for Evaluating AI-Generated Evidence* released in April 2026 by Judge Grimm, Professor Grossman and Mr. Brady is organized the way we think a Note should be, by the question the court is answering rather than by the technology in front of it. We will draft and supply a provenance branch for it, as a contribution to their instrument rather than a proposal for our own, if its authors want one.

An assurance tier defined by what a court needs. Conformance programs for provenance signing are presently designed around what a distribution platform needs. Nothing in them is addressed to what a court needs, and the sixth finding of the security analysis in the agenda book, that the conformance program "does not define security requirements for conforming validators" (page 141), is consistent with that observation. If that gap is not closed, a trial judge will improvise a standard under time pressure in a case that none of us controls, and it will be cited. Encypher is arguing inside the relevant conformance process for a tier with published criteria, independent assessment, and a documented revocation and incident-response process. We mention it because it constrains vendors including us, and because it is the kind of work that belongs outside the Rules rather than inside them.

ENCYPHER

We also note, in the interest of pointing the Committee away from ourselves, that the President's AI Action Plan of July 23, 2025 directs the National Institute of Standards and Technology to consider developing its Guardians of Forensic Evidence deepfake evaluation program "into a formal guideline and a companion voluntary forensic benchmark" (November 2025 agenda book, page 151), and that the lead of that initiative is among the panelists the Committee will hear on October 15 (page 155). We think that is the correct home for anything resembling a technical specification, and that the Committee is right to keep such material out of the Rules. Our recommendations are drafted so that nothing in them requires the Committee to adopt, evaluate, or refer to any standard.

Respectfully submitted,

Eddan Katz
Head of Policy and Government Affairs
Encypher Corporation
2093 Philadelphia Pike #6262, Claymont, Delaware 19703
eddan.katz@encypher.com

ENCYPHER

Annex A. Consolidated proposed rule text

Additions to the working draft reproduced at pages 150 through 151 of the October 15, 2026 agenda book are shown underlined and deletions struck through. The draft is otherwise unchanged.

> **Rule 901(c). Potentially Fabricated Evidence Created by Generative Artificial Intelligence**
>
> **(1) Showing Required to Warrant an Inquiry into Fabrication.** If a party challenges the authenticity of an item of evidence on the ground that it has been fabricated, in whole or in part, by generative artificial intelligence, the party must present evidence sufficient to support a finding of fabrication to warrant an inquiry by the court.
>
> **(2) Showing Then Required by the Proponent.** If the opponent makes that showing, the item of evidence is admissible only if the proponent demonstrates to the court that it is more likely than not authentic.
>
> **(3) Applicability.** This rule applies to items offered under either Rule 901 or 902.
>
> **(4) Notice.** Unless the court orders otherwise, a party challenging an item of evidence under this subdivision (c) must provide reasonable pretrial notice of the intent to present evidence of fabrication, so that all other parties have a reasonable opportunity to respond to that evidence.
>
> **(5) Definitions.** In this rule, "generative artificial intelligence" means a computer system that ~~emulates the structure and characteristics of input data in order to create derived synthetic content, including images, video, audio, text, and other digital content~~ creates content, or that alters existing content, so as to depict a person, object, place, or event that the content does not record.
>
> **(6) Absence of Origin Data.** The absence of data describing the origin or history of an item does not by itself satisfy paragraph (1).

Note on numbering. Paragraph (6) is numbered on the assumption that paragraph (5) survives. If the Committee were to remove the definition from the text, the new provision becomes (5). Its placement matters less than its location in the text rather than the Note.

ENCYPHER

Annex B. Consolidated proposed Committee Note language

The following are offered as insertions into the draft Note reproduced at pages 150 through 152 of the October 15, 2026 agenda book, in the order in which they would appear. Line references are to that draft.

Replacing the paragraph the Reporter proposes to insert as the third paragraph of the Note (pages 152-153).

> Under this rule, the opponent's objection must be grounded in facts raising a genuine concern about the offered item. The showing must be specific to the item challenged: it may rest on a gap or break in the item's chain of custody, on an inconsistency between the item and any record of its own origin or modification or between such a record and other evidence, or on a characteristic of the item itself that is inconsistent with the manner of creation the proponent asserts. A general demonstration that fabrication of such items is possible does not satisfy paragraph (1), nor does the absence of a record of origin by itself. The facts pertinent to the inquiry are likely to change as technology develops, but the following is a non-exhaustive list of facts that would be pertinent to the opponent's burden of going forward:
>
> - metadata that is anomalous, internally inconsistent, or inconsistent with the manner of creation the proponent asserts;
> - inconsistent timestamps or creation dates;
> - unexplained edits or compression artifacts;
> - the proponent's inability to produce the native file when a derivative version is offered, where the proponent's own account indicates that a native file existed;
> - gaps in chain of custody that would raise an opportunity for tampering;
> - a result of a process used to verify a record of the item's origin or modification that indicates alteration of the item or of the record;
> - an anomaly in the path by which the item was acquired;
> - glitches such as unusual audio patterns and problems of synchronization;
> - known use of generative AI tools by the proponent in connection with the item or the case; and
> - independent evidence that is inconsistent with the content of the item.

After line 36 (the nature of the paragraph (2) determination).

> A determination under paragraph (2) is a preliminary question for the court. It does not alter the allocation of the burden of persuasion on any issue in the case, and it does not preclude a party from presenting evidence bearing on authenticity to the trier of fact after the item is admitted.

After line 46 (evidence relevant to paragraph (2)).

> Evidence relevant to a determination under paragraph (2) may include a record of the item's origin or modification that is cryptographically bound to the item, together with the result of a process used to verify that record and evidence of custody, source, device, and workflow. In evaluating such a record the court should consider what portion of the item the binding covers; the results of signature and credential validation, including whether the credential had been revoked or had expired and how the verifying process treated that status; the trust policy applied and disclosed; the identity and security of the system that generated the record; the source and verification of any time stamp; the identity, version and configuration of the process used to verify the record; links to source material;

ENCYPHER

> preservation; and whether the result can be reproduced by another party. A favorable technical result does not, without more, establish the truth of an assertion, the identity of a human author, or the occurrence of a depicted event.
>
> Independent signals bearing on the origin of an item may conflict, and different tools applied to the same record may report different results. A conflict among such signals or results is a fact for the court to resolve on the evidence under paragraph (2). Neither the presence nor the absence of any single signal is controlling.

Replacing or supplementing lines 64 through 66 (Rule 902).

> A certification under Rule 902(11), (12), (13), or (14) does not preclude an inquiry under this subdivision. In determining a question under paragraph (2), the court may consider a certified record of the results of a process used to verify data describing the origin or history of the item, together with the other evidence bearing on authenticity.

Replacing lines 84 through 85 (attribution of the definition). See Annex C for the alternatives. In substance, the Note should not rest the rule's vocabulary on Executive Order 14110, which Executive Order 14148 revoked on January 20, 2025.

For the Rule 707 Committee Note (Recommendation 9), revising the bullet at lines 69-70.

> Reviewing information about what records the system keeps, what is deleted, and whether the output can be reproduced or audited, distinguishing records retained by the system or its operator from records that accompany the output and are available to any party in possession of it.

For the Rule 707 Committee Note (Recommendation 10), after line 60.

> In applying Rule 702(a)-(d) to machine-generated evidence under Rule 707, courts may consider threshold integrity factors necessary for meaningful adversarial testing, including whether the proponent can provide tamper-evident records sufficient to detect missing or altered inputs; identify the system, model version, and execution context that generated the output; and permit independent verification of completeness and provenance. Where such records accompany the output and can be verified by any party in possession of it, the integrity facts they establish may be shown by a certification under Rule 902(13) or (14) rather than by live testimony. A certification of that kind does not establish the reliability of the system's analysis.

For the Rule 707 Committee Note (Recommendation 12).

> Nothing in this rule affects the application of the Confrontation Clause. In a criminal case, the sufficiency of a certification to establish a foundation remains subject to the defendant's right to be confronted with the witnesses against him.

For the Rule 104 Committee Note (Recommendation 11(b)), revising the citation at lines 21-24.

> ... See Rules 104(b), 602, 901(a) and 1008.

For the Rule 104 Committee Note (Recommendation 11(a)), revising the sentence proposed at page 196.

> Nothing in this amendment is intended to affect the admissibility requirements under Rules 602 or 901(a). These are now examples of rules "otherwise provided." See Rule 104(a). [The Committee should state here whether a contested question of authenticity concerning an item offered under Rule 902 is governed by Rule 104(a) or by the standard of Rule 901(a).]

For any future Committee Note to Rule 902(14) (Recommendation 11(c)).

ENCYPHER

> A process of digital identification within the meaning of this rule includes verification of a record of an item's origin or modification that is cryptographically bound to the item. A certification under this rule establishes only the identity of the copy and the technical results described. It does not establish the truth or completeness of any assertion, the identity of a human author, the reliability of machine-produced information, or admissibility under any other rule.

ENCYPHER

Annex C. Statutory and comparative definitions, and three drafting architectures

This annex responds to the statement in the Reporter's April memorandum that "[t]here is no definition of generative artificial intelligence that I could find in a federal statute" (May 2026 agenda book, page 220), and to panel questions 5 and 6. We have checked and believe that statement is correct as of the date of this submission. What follows is the material we would put in its place. Statutory quotations have been checked against the codified text or, for Regulation (EU) 2024/1689, the Official Journal; Part 4 notes the two respects in which we rely on secondary or reproduced text.

1. The one federal statutory definition in this area is of "artificial intelligence"

The definition used in Rule 707(e) originates in the National Artificial Intelligence Initiative Act of 2020, enacted as Division E of the National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, and is codified at **15 U.S.C. § 9401(3)**. The Rule 707 Note cites the public law but not the Code (Note lines 110-111). We recommend adding the Code citation, both because it is more usable and because it makes clear that the definition rests on a statute rather than an executive instrument. The New York court rule reproduced at page 173 adopts the same statutory language.

The statute's further elaboration, which the Rule 707 Note already reproduces, is that artificial intelligence systems "use machine and human-based inputs to - (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action."

There is no corresponding federal statutory definition of "generative artificial intelligence."

2. The three architectures now in use

Architecture A: process and training data. This is the family to which the current draft belongs. Executive Order 14110, § 3(p), defined generative AI as "the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content." California has enacted a close variant: **Cal. Gov't Code § 11549.64(b)** defines generative artificial intelligence as an artificial intelligence system "that can generate derived synthetic content, including text, images, video, and audio that emulates the structure and characteristics of the system's training data," and the same definition appears in the California AI Transparency Act (Cal. Bus. & Prof. Code § 22757.1(g)). A parallel formulation appears in California's training-data transparency statute, keyed to "the artificial intelligence's training data" (Cal. Civ. Code § 3110(c)). The Connecticut Practice Book provision reproduced at page 135 is a degenerate member of the same family: "any machine-learning model/application that can create original content ... in response to a prompt or request," followed by a list of named products.

Assessment. The Committee should know that the substance of the draft definition survives in enacted state law even though Executive Order 14110 was revoked. If the Committee wishes to keep this formulation, the simplest correction is to re-attribute it: cite the California provision rather than the rescinded order. That cures the attribution problem at no substantive cost. It does not cure the three substantive defects discussed in Recommendation 3: the formulation is not comprehensible on its face, it is architecture-descriptive, and its reach over generative alteration of authentic content is uncertain. The Connecticut variant shows where product- and architecture-naming leads.

Architecture B: capability. 15 U.S.C. § 9401(3), above. Appropriately broad for Rule 707, which regulates analytical output. Far too broad for Rule 901(c), which regulates depiction.

Architecture C: output and resemblance. Regulation (EU) 2024/1689, article 3(60), defines a deep fake as "AI-generated or manipulated image, audio or video content that resembles existing persons, objects, places, entities or events and would falsely appear to a person to be authentic or truthful."

Assessment. Two features are instructive. "AI-generated or manipulated" reaches both wholly synthetic output and alteration of authentic content, which is the gap in the current draft. But the trailing clause, "would falsely appear to a person to be authentic or truthful," is a reasonable-observer test, and the Reporter's objection to the TAKE IT DOWN Act's equivalent applies to it with equal force: such a test is out of place in a rule under which the court determines whether an item is more likely than not authentic. The European definition is also limited to image, audio and video, which the Reporter has confirmed the draft rule is not (page 158, panel question 4).

The **TAKE IT DOWN Act**, Pub. L. No. 119-12, § 2, 139 Stat. 55 (May 19, 2025), codified at 47 U.S.C. § 223(h)(1)(B), belongs in this family. It defines "digital forgery" as:

> any intimate visual depiction of an identifiable individual created through the use of software, machine learning, artificial intelligence, or any other computer-generated or technological means, including by adapting, modifying, manipulating, or altering an authentic visual depiction, that, when viewed as a whole by a reasonable person, is indistinguishable from an authentic visual depiction of the individual.

Assessment. The Reporter is right to reject this definition as a whole, for the reason he gives and for the further reason that it is confined to intimate visual depictions of identifiable individuals. But the clause **"including by adapting, modifying, manipulating, or altering an authentic visual depiction"** is precisely the coverage the current draft lacks, and it is available from an enacted federal statute. We recommend borrowing it in substance while declining the reasonable-person test.

Louisiana's deepfake provision, La. R.S. 14:73.13, which the Reporter cited in November 2025, also belongs in this family and should be included in the verified compilation.

3. The architecture we recommend: output relative to the event

> a computer system that creates content, or that alters existing content, so as to depict a person, object, place, or event that the content does not record.

Why. It covers synthesis and generative alteration in one clause, which no single existing definition does without importing a reasonable-observer test. It excludes enhancement that clarifies what was recorded, which is the exclusion the Reporter wants for sharpening filters and speech-to-text transcription, while capturing enhancement that invents detail, which is what the courts found in *State v. Puloka* and *Linkugel v. City of Los Angeles*. It contains no reasonable-person test, so the determination remains the court's, consistently with paragraph (2). It applies to any item of evidence rather than to audiovisual material only, consistently with the Reporter's confirmation that the rule "covers all items of evidence" (page 158). And it describes a relationship between an output and the world rather than a system architecture, so it does not require amendment when architectures change.

Limits we acknowledge. The formulation places weight on "record," which will need to bear the distinction between a camera capturing an event and a model producing an image of one. We think that is the right place for the weight to sit, because it is the distinction the rule is about. We would welcome the chance to

ENCYPHER

test it against harder cases, including composite items with both recorded and generated regions, and will do so in writing after the meeting if the Reporter would find it useful.

4. A note on sources

California. The California AI Transparency Act has been amended since the version reproduced at agenda book pages 136-137. Where this submission relies on the Act's current requirements (Section IV), it cites the codified sections, Cal. Bus. & Prof. Code §§ 22757.1-22757.3.3; where it discusses the Reporter's note on the Act, it cites the agenda book.

Foreign decisions. The New South Wales and Ontario decisions described in Section III are cited as quoted in the Victorian Law Reform Commission's report, *Artificial Intelligence in Victoria's Courts and Tribunals* (October 2025), paras 5.121-5.122. We have not obtained the primary judgments.

ENCYPHER

Annex D. Authorities cited

Committee materials

- Advisory Committee on Evidence Rules, agenda book for the meeting of October 15, 2026 (221 pp.), including: draft minutes of the meeting of May 7, 2026, pp. 13-42; draft minutes of the June 2026 meeting of the Standing Committee, pp. 43-85; Reporter's memorandum, *Deepfakes and Draft Rule 901(c)*, September 15, 2026, pp. 133-160; Reporter's memorandum, *Artificial Intelligence and Proposed Rule 707*, September 15, 2026, pp. 162-182; Academic Consultant's memorandum, *Proposed Amendment to Rule 902(1)*, September 18, 2026, pp. 184-185; Reporter's memorandum, *Possible Amendments to Rule 104(a) and (b)*, September 15, 2026, pp. 187-196.
- Agenda book for the meeting of May 7, 2026 (355 pp.), including: Reporter's memorandum on Artificial Intelligence and Proposed Rule 707, April 5, 2026, pp. 116-179; letter of the Department of Justice, February 13, 2026, pp. 180-192; written testimony of Nicole Owens, Executive Director, Federal Defender Services of Idaho, January 15, 2026, pp. 193-196; Reporter's memorandum on Deepfakes and Draft Rule 901(c), April 11, 2026, pp. 198-225; Federal Judicial Center, *Summary of Rule 901 Survey*, March 25, 2026 (Ashley Walters, Meghan Dunn, Timothy Lau), pp. 226-232; Reporter's memorandum on Possible Amendment to Rule 104(a) and (b), April 1, 2026, pp. 333-355.
- Agenda book for the meeting of November 5, 2025 (317 pp.), pp. 102-175.
- Report of the Advisory Committee on Evidence Rules, December 1, 2025.
- Proposed Amendments to the Federal Rules of Evidence, Docket No. USC-RULES-EV-2026-0036 (comment period August 14, 2026 to February 15, 2027).

Cases

- *Adams v. Anderson*, 2026 WL 1501021 (Ky. Ct. App. May 29, 2026).
- *Al-Qarqani v. Chevron Corp.*, No. 19-17074, 2023 U.S. App. LEXIS 16081 (9th Cir. June 26, 2023).
- *Bullcoming v. New Mexico*, 564 U.S. 647 (2011).
- *Burnley v. Valentin*, 2026 WL 767145 (E.D. Va.).
- *Higgins v Commissioner of Police, NSW Police Force* [2024] NSWCATAD 175 (as reported by the Victorian Law Reform Commission).
- *Huang v. Tesla, Inc.*, No. 19CV346663 (Cal. Super. Ct. Santa Clara Cnty. Apr. 27, 2023).
- *Linkugel v. City of Los Angeles*, 2026 WL 1430790 (Cal. Ct. App.).
- *Matter of M.S.*, 2026 N.Y. Slip Op. 00825, 2026 WL 436359 (N.Y. Feb. 17, 2026).
- *Melendez-Diaz v. Massachusetts*, 557 U.S. 305 (2009).
- *Mendones v. Cushman & Wakefield, Inc.*, No. 23CV028772, 2025 WL 2613764 (Cal. Super. Ct. Alameda Cnty. Sept. 9, 2025).
- *People v. Goldsmith*, 59 Cal. 4th 258 (2014).
- *People v. Wakefield*, 38 N.Y.3d 367 (2022).
- *R v Gallerno* [2025] ONSC 236 (as reported by the Victorian Law Reform Commission).
- *Smith v. Arizona*, 602 U.S. 779 (2024).
- *State v. Puloka*, No. 21-1-04851-2 KNT (Wash. Super. Ct. King Cnty. Mar. 29, 2024).
- *State v. Sample*, 468 Md. 560 (2020).
- *United States v. Andrews*, 172 F.4th 969 (7th Cir. 2026).

ENCYPHER

- *United States v. Belt*, 181 F.4th 931 (8th Cir. 2026).
- *United States v. Dunnican*, 961 F.3d 859 (6th Cir. 2020).
- *United States v. Gogic*, 2025 WL 3042350 (E.D.N.Y.).
- *United States v. Hernandez*, No. 3:25-cr-2384 (W.D. Tex.), sentencing transcript, ECF No. 28.
- *United States v. Martinez*, 2026 WL 544719 (C.D. Cal.).

Statutes, rules and instruments

- 15 U.S.C. § 9401(3) (National Artificial Intelligence Initiative Act of 2020, Pub. L. No. 116-283, div. E).
- 47 U.S.C. § 223(h)(1)(B) (TAKE IT DOWN Act, Pub. L. No. 119-12, § 2, 139 Stat. 55 (2025)).
- Executive Order 14110, § 3(p) (Oct. 30, 2023), revoked by Executive Order 14148 (Jan. 20, 2025).
- *America's AI Action Plan* (The White House, July 23, 2025).
- Cal. Bus. & Prof. Code §§ 22757.1-22757.3.3 (California AI Transparency Act, as amended); §§ 22757.2-22757.3 as reproduced at agenda book pages 136-137.
- Cal. Civ. Code § 3110(c).
- Cal. Gov't Code § 11549.64(b).
- Connecticut Practice Book § 4-9, as reproduced at agenda book page 135.
- La. Code Civ. Proc. arts. 371(C), 1551(A)(5) (Acts 2025, No. 250 (H.B. 178), eff. Aug. 1, 2025); La. R.S. 14:73.13.
- 22 N.Y.C.R.R. Part 161, § 161.2, as reproduced at agenda book page 173.
- 12 V.S.A. § 1913 (enacted 2016, as amended through 2024).
- Regulation (EU) 2024/1689 (Artificial Intelligence Act), arts. 3(60), 50(2), 50(4), OJ L, 12.7.2024.
- Fed. R. Civ. P. 37(e).

Technical and scholarly sources

- Craig Ball, practitioner guidance on managing deepfake problems (August 2026), as reproduced at agenda book pages 137-139.
- Jeffrey Bellin, draft article on AI and evidence, as reproduced at agenda book pages 163-167.
- Maria Berkenkotter and Lino Lipinsky de Orlov, *Judges On AI: How Courts Can Survive The Tech Revolution*, Law360, as reproduced at agenda book page 145.
- Aleeza Furman, *As Deepfake Evidence Spreads, Rulemaking Efforts Stay Stuck in Development*, Law.com (May 7, 2026), as reproduced at agenda book pages 147-149.
- Richard Gilman, *The Legal Implications of Deepfakes: What Courts Are Starting to Say*, Lawbit Digest, as reproduced at agenda book pages 145-147.
- Golaszewski, Krawetz, Sherman et al., *Verifying Provenance of Digital Media: Security Analysis of C2PA and its Implementation*, IACR ePrint 2026/804, summarized as *Verifying Provenance of Digital Media: Why the C2PA Specifications Fall Short* at agenda book pages 139-142.
- Hon. Paul W. Grimm (ret.), Maura R. Grossman and Kevin F. Brady, *Decision Tree for Evaluating AI-Generated Evidence*, 27 Sedona Conf. J. (forthcoming 2026) (preprint April 2026).
- Maura R. Grossman and Hon. Paul W. Grimm, *Judicial Approaches to Acknowledged and Unacknowledged AI-Generated Evidence*, 26 Colum. Sci. & Tech. L. Rev. (2025).
- Edward Imwinkelried, *The Challenge that the Advent of Artificial Intelligence (AI) Tools Poses to the Procedures for Determining the Existence of the Preliminary Facts that Condition the Admissibility of Items of Evidence*, 108 Marq. L. Rev. 621 (2025).

ENCYPHER

- Kassis and Hengartner, *UnMarker: A Universal Attack on Defensive Image Watermarking*, IEEE Symposium on Security and Privacy (2025).
- William Matthewman, *The Coming Use and Misuse of Artificial Intelligence in the Courtroom: A Judicial Perspective and Proposal*, 30 J. Tech. L. & Pol'y 139 (2026), as reproduced at agenda book pages 142-145.
- National Institute of Standards and Technology, *NIST AI 100-4: Reducing Risks Posed by Synthetic Content* (November 2024).
- National Security Agency, cybersecurity information sheet on content credentials (January 2025).
- Alexander Nemecek, Hengzhi He, Guang Cheng and Erman Ayday, *Authenticated Contradictions from Desynchronized Provenance and Watermarking*, arXiv:2603.02378 (March 2, 2026).
- Gregory D. Schwartz, *When Disciplines Disagree: The Admissibility of Computer-Generated Forensic Evidence in the Criminal Justice System*, 72 UCLA L. Rev. Disc. 174 (2024), as reproduced at agenda book pages 167-168.
- Victorian Law Reform Commission, *Artificial Intelligence in Victoria's Courts and Tribunals* (Report, October 2025).
- Weitel, *Defining Artificial Intelligence*, 71 Wayne L. Rev. 363 (2026), as reproduced at agenda book pages 169-172.
- Benjamin Welton, *Meaningful Machine Confrontation*, 76 Stan. L. Rev. 845 (2025), as reproduced at agenda book pages 168-169.
- Jessica Young et al., *Media Integrity and Authentication: Status, Directions, and Futures* (Microsoft Research, arXiv:2602.18681, February 2026).

ENCYPHER

Encypher Corporation: Suggestion to the Advisory Committee on Evidence Rules

Draft Rule 901(c), revised Rule 707, and the published amendments to Rules 104 and 902(1). Summary, October 9, 2026.

Who we are and what we disclose. Encypher builds content provenance infrastructure on the C2PA standards. It operates a claim-signing certificate authority and a time-stamp authority whose credentials appear on the C2PA trust lists, and co-chairs the C2PA task force that wrote the specification's text provisions. Two of the six findings in the security analysis at agenda book pages 139-142 concern functions we perform.

What we do not ask for. Earlier publication of Rule 901(c); any presumption, safe harbor or weight for a provenance record; the naming of any standard, vendor or product in any rule; any inference from the absence of a provenance record; any suggestion that a favorable technical result proves what an item depicts or who made it.

Four recommendations run against our commercial interest: 1, 5, 7 and 12.

The twelve recommendations, in one line each.

1. Add Rule 901(c)(6): "The absence of data describing the origin or history of an item does not by itself satisfy paragraph (1)." Absence is uninformative because ordinary handling strips records, and is systematically correlated with the criminal file. Where the proponent's own account implies a record should exist, including on capture devices California will require to embed provenance from 2028, the inference remains available.
2. Adopt the Reporter's factor paragraph with a functional chapeau (custody breaks, inconsistencies with records or other evidence, characteristics inconsistent with the claimed manner of creation) and four edits, chiefly separating "missing" from "anomalous" metadata.
3. Define generative AI by what the output depicts: a system that "creates content, or that alters existing content, so as to depict a person, object, place, or event that the content does not record." Re-attribute the definition away from revoked Executive Order 14110.
4. Keep the notice provision as adopted, without a good-cause exception.
5. A conflict among signals or verification tools is a fact for the court, not a verdict.
6. Tell courts what a favorable technical result does and does not establish, with factors that track each finding of the security analysis.
7. A Rule 902(11)-(14) certification does not end a Rule 901(c) inquiry; a certified verification record may be considered on the shifted question.
8. Confirm that paragraph (2) does not alter the burden of persuasion.
9. Keep the three new Rule 707 analysis bullets; distinguish records kept by the system from records that travel with the output.

ENCYPHER

10. Permit the integrity foundation for acknowledged machine output (which system, whether altered, whether reproducible) to be shown by certification, never the model's reliability, and only together with Recommendation 12.
11. Resolve in the Rule 104 Note whether contested Rule 902 items fall under amended Rule 104(a); cite "Rule 901(a)"; clarify Rule 902(14)'s reach to cryptographic records of origin.
12. Confirm that any certification route is subject to the Confrontation Clause.

Offers. A verified compilation of statutory definitions; published filing-system survival testing, whatever it shows; a free verifier that returns full results rather than a badge; a technical annex within thirty days of the meeting; a provenance branch for the Sedona decision tree, if its authors want one; and advocacy for a court-facing assurance tier in the C2PA conformance process.

Contact. Eddan Katz, Head of Policy and Government Affairs, eddan.katz@encypher.com.