Machine Translated by Google

RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

Reconfiguring evidence discovery: the proportionality standard and defensibility in the era of generative e-discovery.

Reconfiguring discovery: the standard of proportionality and defensibility in the era of generative e-discovery

Reconfiguring test discovery: the standard of proportionality and defense in the era of generative e-discovery

Vinicius Feitosa Veras

ABSTRACT

The massive volume of electronically stored information (ESI), driven by the widespread adoption of corporate collaboration tools and, more recently, by generative artificial intelligence (Generative AI) models, is straining traditional evidence discovery (e-discovery) structures in the American civil procedure. This article aims to analyze how the proportionality standard enshrined in Rule 26(b)(1) of the Federal Rules of Civil Procedure (FRCP) is being reconfigured in the face of this new scenario, with emphasis on the phenomenon of "Shadow AI"—the ungoverned corporate use of generative AI tools by employees—and the consequent need for next-generation ESI protocols.

Through qualitative, exploratory, legal-doctrinal, and documentary research, paradigmatic judicial precedents are examined, such as Da Silva Moore v. Publicis Groupe (2012) and Mata v. Avianca, Inc. (2023), as well as guidelines emanating from the Sedona Conference, to demonstrate the paradigm shift from Technology-Assisted Review (TAR) to generative AI. The results indicate that procedural defensibility—that is, the ability of a discovery process to withstand judicial challenges—increasingly depends on traceability, auditability, and algorithmic governance mechanisms that have not yet received uniform treatment in jurisprudence or procedural law. The conclusion is that ESI protocols must expressly incorporate prompts, generated outputs, and metadata of interaction with AI systems, otherwise both the proportionality and integrity of the duty to preserve evidence will be compromised.

Keywords: E-discovery. Generative Artificial Intelligence. Proportionality. Shadow AI. Electronically Stored Information (ESI).

ABSTRACT The

massive volume of Electronically Stored Information (ESI) generated by corporate collaboration tools and, more recently, by generative artificial intelligence (Generative AI) models is straining the traditional structures of e-discovery in United States civil litigation. This article aims to analyze how the proportionality standard enshrined in Rule 26(b)(1) of the Federal Rules of Civil Procedure (FRCP) is being reconfigured in light of this new landscape, with emphasis on the phenomenon of "Shadow AI" — the ungoverned corporate use of generative AI tools by employees — and the consequent need for next-generation ESI protocols. Through qualitative, exploratory, doctrinal and documentary legal research, the study examines paradigmatic judicial precedents, such as Da Silva Moore v. Publicis Groupe (2012) and Mata v. Avianca, Inc. (2023), as well as guidance issued by The Sedona Conference, to demonstrate the paradigm shift from Technology-Assisted Review (TAR) to generative AI. The findings indicate that procedural defensibility — that is, the capacity of a discovery process to withstand judicial challenge — increasingly depends on traceability, auditability and algorithmic governance mechanisms that still lack uniform treatment in case law or procedural legislation. The article concludes that ESI protocols must expressly incorporate prompts, generated outputs and AI-interaction metadata, under penalty of compromising both proportionality and the integrity of the duty to preserve evidence.

Keywords: E-discovery. Generative Artificial Intelligence. Proportionality. Shadow AI.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

Electronically Stored Information (ESI).

1 INTRODUCTION

Contemporary civil litigation in the United States faces a structural challenge: the volume of data potentially relevant to a legal dispute is growing at a much faster rate than the ability of the parties, lawyers and courts to process it in a proportionate manner and defensible. This phenomenon, already noticeable since the consolidation of the Electronically Stored concept. Information Security Information (ESI) intensified in the 2006 reforms to the Federal Rules of Civil Procedure (FRCP). with the spread of corporate collaboration tools — such as messaging platforms Instant messaging, cloud-based work environments, and project management systems — and more. recently, with the everyday incorporation of generative artificial intelligence (AI) models Generative) to the work routines of employees and organizations.
Generative AI can produce not only new records—prompts, responses, and metadata. Conversational — as it is also frequently adopted by employees without regard to governance. Formal information technology in companies, a phenomenon known as "Shadow AI".
Research in the information security sector indicates that many organizations do not have complete visibility into the use of generative AI tools by your employees and that A significant portion of employees share sensitive data with these tools without Corporate authorization (CYBERHAVEN, 2025; IBM, 2025; NEURALTRUST, 2026). This scenario This calls into question two structural pillars of the American civil procedure: the standard of proportionality of Rule 26(b)(1) of the FRCP, which presupposes the possibility of mapping and delimiting rationally the universe of ESI potentially subject to discovery, and the duty of preservation of evidence, the violation of which may give rise to sanctions under Rule 37(e).

1.1 Research Problem

Given this scenario, the following research problem is formulated: to what extent does the pattern of proportionality established by Rule 26(b)(1) of the FRCP and traditional ESI protocols remain capable of ensuring processes for discovering defensible evidence in the face of the explosion of volume, algorithmic opacity, and fragmented governance resulting from corporate use — including unauthorized generative artificial intelligence tools?

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

1.2 Objectives

The overall goal of this article is to analyze how the proportionality pattern in e-discovery is...
being reconfigured by the rise of generative artificial intelligence, identifying the main
Risks to procedural defense and the elements that should comprise adequate ESI protocols.
to this new technological context.
Specifically, the aim is to: (i) describe the conceptual evolution of ESI and the standard of
proportionality under Rule 26(b)(1) of the FRCP; (ii) characterize the phenomenon of Shadow AI and its
impacts on corporate data governance and on the duty to preserve evidence; (iii)
to examine paradigmatic judicial precedents — notably Da Silva Moore v. Publicis Groupe
(2012) and Mata v. Avianca, Inc. (2023) — which highlight the transition from assisted review to
technology (TAR) for generative AI in litigation; and (iv) propose structuring elements of
next-generation ESI protocols, compatible with proportionality requirements and
procedural defensibility.

1.3 Justification

The relevance of the topic stems from three sets of reasons. Firstly, practical reasons: offices.
Law firms, legal departments, and courts already encounter disputes in their daily legal practice.
regarding the discovery of interaction records with generative AI tools, without them existing
Established treatment patterns. Secondly, doctrinal reasons: the legal literature.
specialized in e-discovery, although extensive in terms of technology-assisted review (TAR) and
Traditional proportionality still offers an incipient and fragmented approach to AI.
generative, especially regarding Shadow AI as an autonomous source of evidentiary risk. Thirdly
Location, reasons for relevance: recent institutional initiatives — such as the working groups of
Sedona Conference dedicated to artificial intelligence (THE SEDONA CONFERENCE, 2026) —
They demonstrate that the American justice system itself is in an active process of development.
of guidelines, which makes a critical and systematic analysis of this topic particularly timely.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

2. THEORETICAL FRAMEWORK

2.1 ESI and the traditional architecture of electronic discovery

The concept of Electronically Stored Information (ESI) was formally incorporated into the Federal Rules. of Civil Procedure by the 2006 amendments, which deliberately broadened the scope of Rule 34. to encompass information stored electronically. The technologically open formulation It sought to prevent the rules from becoming obsolete in the face of new forms of storage. This The regulatory option allows, in principle, access to records of interaction with intelligence systems. artificial generative reproduction, provided that the requirements of relevance, proportionality and ownership are met, custody or control applicable to the discovery. The specific extent of these obligations, however, must be examined on a case-by-case basis, especially when the data is under vendor infrastructure. third parties (UNITED STATES, 2015).

The traditional architecture of electronic discovery rests on three pillars: (i) the obligation to preservation, which arises from the moment the litigation is reasonably foreseeable; (ii) a obligation to produce relevant and non-privileged documents; and (iii) the possibility of judicial control over the scope, method, and cost of discovery, through the standard of proportionality. These pillars were conceived in a context in which the ESI universe, Although voluminous, it was relatively stable and identifiable — emails, spreadsheets, databases. Corporate. The rise of generative AI conversational tools introduces a fourth Element of complexity: data that are simultaneously the product of a probabilistic process. and potentially non-reproducible, which challenges established hermeneutical assumptions about the which constitutes a "complete" or "authentic" document for evidentiary purposes.

2.2 The proportionality standard of Rule 26(b)(1) of the FRCP

Rule 26(b)(1) of the FRCP delimits the general scope of evidence discovery in federal civil proceedings. North American, authorizing the parties to obtain discovery on any non-privileged matter. that is relevant to the claim or defense of any party and proportionate to the needs of the case. The text of the rule, in the version in force since the 2015 amendment, is explicit regarding the factors that The following elements comprise the proportionality assessment:

> Parties may obtain discovery regarding any nonprivileged matter that is relevant to any party's claim or defense and proportional to the needs of the case, considering the importance

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

> of the issues at stake in the action, the amount in controversy, the parties' relative access to relevant information, the parties' resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit (UNITED STATES, [2015], emphasis added).

The 2015 amendment did not create the concept of proportionality, which was already present in the Rule.
26(b)(2)(C), but repositioned it at the core of the definition of the scope of the discovery. The change made
Proportionality is a constitutive element of the discovery itself, and not merely a limitation.
subsidiary. This does not mean, however, that the requesting party has received the exclusive burden of
demonstrate all factors of proportionality: according to the Advisory Committee's note, the parties and
the court shares the responsibility of considering these factors, and a generic objection from
The fact that the party is required is not sufficient to refuse production (UNITED STATES, 2015).
The proportionality assessment operates as a global cost-benefit calculation, which weighs the...
The importance of the issues in dispute, the value of the case, the parties' relative access to information,
available resources, the importance of discovery for dispute resolution, and the relationship between the burden of proof.
or the cost of the intended discovery and its likely benefit. None of these factors, in isolation,
decisive (LITIGATION AND TRIAL apud GLOBAL RELAY, 2024). This normative design
However, it assumes that it is possible to estimate, with reasonable accuracy, the volume, location, and
cost of access to the relevant ESI universe — an assumption that is becoming increasingly fragile.
as an increasing portion of corporate communication takes place in generative AI tools
not inventoried by the information technology area.

2.3 From TAR to generative AI: the technological genealogy of document review

The technological response traditionally employed to reconcile the growth in the volume of ESI.
The standard of proportionality was the technology-assisted review (Technology-Assisted
Review (TAR), also referred to as "predictive coding" or computer-assisted review. The milestone
The foundational jurisprudential precedent for the judicial recognition of TAR is the decision issued by the Magistrate.
Andrew J. Peck, of the Southern District Court of New York, in the case Da Silva Moore v. Publicis
Groupe & MSL Group, 287 FRD 182 (SDNY 2012). This was an action for discrimination of
A lawsuit filed by former employees against a communications company, involving a universe of...
approximately three million potentially relevant electronic documents (DA SILVA
MOORE V. PUBLICIS GROUPE, 2012; EVERLAW, 2024).
Judge Peck acknowledged that manually reviewing a volume of that magnitude would be economically...
costly and subject to inconsistencies between reviewers, concluding that algorithm-assisted review

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

of machine learning, when properly validated by sampling and iterative training,
This could constitute an acceptable method for locating relevant ESI in appropriate cases (DA SILVA
(MOORE V. PUBLICIS GROUPE, 2012). The magistrate's decision was subsequently adopted by
district judge, who rejected the objections presented by the authors. Three years later, in Rio Tinto
PLC v. Vale SA, Judge Peck himself observed that the judicial acceptance of the TAR had already become
widely established in cases with a large volume of documentation (EVERLAW, 2024).
Generative AI introduces qualitatively different challenges from those faced by ANT. Systems of
TARs typically assign relevance scores or ratings to documents based on...
Examples previously coded by human reviewers; large-scale language models,
In turn, they generate probabilistic responses to instructions in natural language. While processes
TAR (Tactical Action Report) can be evaluated by metrics such as accuracy, recall, and F-measure, the use of...
Generative AI also requires controls focused on the variability of responses and the traceability of...
interactions and the risk of generating plausible but factually incorrect or nonexistent content.
(THE SEDONA CONFERENCE, 2026).

2.4 Shadow AI: The New Frontier of Corporate Risk

The term "Shadow AI" refers to the use of tools by employees or business units to...
artificial intelligence — especially generative intelligence — without formal knowledge, approval, or governance.
from the organization's information technology and information security area (IBM, 2025). The
The phenomenon is analogous, in its logic, to the well-known "Shadow IT," but presents risks.
qualitatively distinct: while shadow IT applications typically store data,
Generative AI tools process, incorporate, and, in certain configurations, can retain...
content submitted by the user as part of their own operating process
(CYBERHAVEN, 2025).
Surveys from the information security sector converge on the relevance of the phenomenon.
Although the percentages vary depending on the methodology used. Gartner survey with leaders.
A cybersecurity survey of 302 organizations identified that many of them suspected or already had
The use of unauthorized generative AI tools by employees has been confirmed.
(NEURALTRUST, 2026). The Cisco AI Readiness Index of 2025, conducted with eight thousand decision-makers.
A study on information technology decision-making in thirty countries highlighted limitations in visibility regarding the
Effective use of AI tools at work (CISCO apud ITBREW, 2025). The cost report of
IBM data breach relating to the period between March 2024 and February 2025 also

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

associated incidents involving Shadow AI with additional security costs (IBM apud ITBREW, 2025). As this is industry-generated evidence, this data should be interpreted with

Caution is advised, and these guidelines do not replace specific empirical research on litigation.

Concrete cases illustrate the risk from the perspective of evidence discovery. Widely publicized reports In the corporate security sector, they describe episodes in which engineers from companies... technology companies inserted proprietary source code and transcripts of internal meetings into tools. Generative AI for public use, leading the organization to urgently review its policy on use of these tools (VECTRA AI apud CLOUD SECURITY ALLIANCE, 2025; CBIZ, 2025).

From an e-discovery perspective, each of these episodes generates a data trail — prompts, Responses, session metadata — which constitutes potentially relevant ESI, but which, by definition, escapes the formal inventories of information systems maintained by the technology area of institution, compromising both the ability to identify and the ability to preserve

This information will be provided when required in court.

2.5 The Sedona Conference and the first efforts to standardize generative AI in the process

The Sedona Conference, a non-profit organization dedicated to improving the law. In civil procedure and e-discovery practices in the United States, it has historically played a role. quasi-normative in the formation of professional consensus on defensible discovery practices electronics, such as the already established Sedona Principles, a widely used doctrinal reference. cited by federal courts since the early 2000s (THE SEDONA CONFERENCE, [n.d.]).

Given the rise of generative AI from 2022 and 2023 onwards, the organization's Working Group 1, responsible for the management and production of electronic documents, formed a specific group. of studies on machine learning and generative AI, culminating in the organization of Thematic conferences dedicated exclusively to artificial intelligence and law, starting in April. 2024 (THE SEDONA CONFERENCE, 2024).

In 2025 and 2026, the organization made progress in developing specific guidelines, including a document. aimed at guiding judges on the use of AI and generative AI in the performance of their duties. judicante (THE SEDONA CONFERENCE, 2025) and a draft, still in the comments phase. public, from a "Primer" on the use of generative AI in evidence discovery, which examines, among other things, other topics, the technical fundamentals of LLMs, the main use cases of technology in litigation and risk mitigation strategies informed by prior experience with TAR. (THE SEDONA CONFERENCE, 2026). The circumstance that this document remains, until the

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

This research, currently in draft form and open to comments, reveals that the legal community
The specialized field has not yet reached a consolidated consensus on defensibility parameters for the
Generative AI in e-discovery, which reinforces the scientific relevance of the present study.

3. TECHNICAL DEVELOPMENT: THE ARCHITECTURE OF E-DISCOVERY GENERATIVE

3.1 Expanding the ESI universe: prompts, generated outputs, and conversational metadata

The adoption of generative AI tools can potentially expand the categories of data.
relevant in a discovery process. According to the architecture and retention policies of each
In this system, these categories may include: (i) prompts entered by users; (ii) generated outputs and
eventually preserved; (iii) available session metadata, such as date, time and version of
model; and (iv) derivative artifacts, such as documents, emails, or procedural documents produced with
Tool support. The subjection of each category to production is not automatic: it depends on
relevance, proportionality, technical availability, and possession, custody, or control by the party.
This potential expansion is compatible with the technologically open formulation of ESI adopted by
FRCP. Still, it should not be assumed that every record of interaction with generative AI is
necessarily stored, accessible, or subject to production. The legal and practical difficulty lies
in determining what data exists, where it is stored, who controls it, and what
Preservation measures are reasonable, especially when interactions occur in tools of
Shadow AI is outside the organization's formal inventory.

3.2 The problem of "reasonable search" in generative AI environments

The "reasonable search" standard, long established in e-case law,
Discovery requires the party obligated to produce evidence to undertake proportionate and good-faith efforts.
faith to identify and collect relevant ESI, without being required to exhaust all and any
a conceivable source of data. This standard presupposes, however, the existence of a reasonably reliable map.
complete access to corporate data sources — an assumption that Shadow AI directly erodes. If the
The legal department of an organization is unaware of what percentage of its employees use tools for
Non-corporate generative AI for drafting communications, analyzing data, or reviewing contracts, makes-
if it is structurally impossible to include these sources within the scope of a reasonable search, even if

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

good faith.

This discrepancy creates a double risk. On the one hand, the defendant may, in good faith, underestimate the... relevant ESI universe, exposing itself to allegations of incomplete production and, in more serious cases, requests for sanctions for failure to fulfill the duty of preservation. On the other hand, the requesting party may seek, through overly broad discovery requests, access to any and all records of interaction with generative AI tools in an organization, under the argument that any One of these tools could, in theory, have been used to address issues related to dispute — a claim that tends to rightfully clash with the very standard of proportionality of Rule 26(b)(1), insofar as the burden of an indiscriminate search over dozens of tools The use of uninventoried AI will hardly be considered proportional to the expected benefit. (LOGIKCULL, 2018).

3.3 Next-generation ESI protocols: proposed technical elements

Traditional e-discovery literature recognizes the importance of the meeting between the parties ("meet and confer", pursuant to Rule 26(f) of the FRCP) for the definition of scope, format and costs from the discovery, including regarding data sources that are difficult to access (JATHEON, 2026). As a proposal normative aspect of this article — and not as an obligation already uniformly established. —, It is argued that ESI protocols negotiated at this stage should consider specific clauses regarding generative AI. structured around four axes:

a) AI data mapping: inventory, updated periodically, of corporately sanctioned generative AI tools and, to the extent possible, evidence of Unsanctioned use, with indication of responsible technicians, associated databases and policies. applicable retention periods;

b) Custody and preservation of prompts and outputs: definition of those responsible for the custody of records. of interaction (prompt and output logs), with retention consistent with the duty of preservation. arising from reasonably foreseeable litigation, including when the data resides on servers of third-party providers of the AI tool;

c) Model versioning: recording the specific version of the AI model. generative engineering used in a given period, insofar as different models — or even Different versions of the same model can produce different outputs from prompts. identical, which is relevant for the reproducibility and interpretation of AI-generated content. in an evidentiary context;

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

d) Traceability and algorithmic chain of custody: when technically
Available and proportionate to the needs of the case, documentation of the data flow from insertion.
from the prompt to the eventual incorporation of the output into a final document, in order to allow the
Reconstructing the role played by generative AI in the production of the artifact.
These elements directly relate to the accumulated experience in TAR protocols, in
What is the level of transparency regarding the training methodology, statistical sampling criteria, and...
Accuracy and recall rates have become best practice to ensure the defensibility of
review process in the event of a judicial challenge (DA SILVA MOORE V. PUBLICIS
(GROUPE, 2012). The essential difference is that, in generative AI, the object of auditing is not only
not a static classifier, but an interactive system whose outputs can vary between sessions, which
requires additional layers of technical documentation.

3.4 Defensibility and the risk of hallucination: Mata v. Avianca as a case study

The case Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (SDNY 2023), became the most significant precedent.
widely cited regarding the risks of generative AI in the American civil process, although not
It's not technically about discovering evidence, but about drafting legal documents.
In this case, lawyers representing the plaintiff in a personal injury lawsuit submitted, before
the Southern District Court of New York, a memorial containing citations to six judicial decisions.
Entirely fictitious, generated by the ChatGPT tool and never verified for their existence.
real (MATA V. AVIANCA, 2023).
Judge P. Kevin Castel, in imposing financial penalties on the lawyers involved, based on...
Rule 11 of the FRCP and the inherent authority of the court did not condemn the use of AI tools.
generative in itself, but the absence of independent human verification of the generated content beforehand
its submission in court (MATA V. AVIANCA, 2023). The episode highlights, in a way
paradigmatically, the phenomenon of "hallucination"—the generation, through language models, of content
syntactically coherent and stylistically plausible, but factually nonexistent — and its relevance.
directly related to the integrity of the judicial process.
From an e-discovery perspective, the precedent offers a relevant analogy, but does not constitute...
decision on eDiscovery. First, it reinforces the desirability of human verification.
regardless of content produced with the support of generative AI before its presentation to
court. Secondly, it shows that reports, memos and internal analyses prepared with
AI support may contain inaccurate claims that, when preserved as corporate records,

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

and relevant to the controversy, they can be integrated into the ESI universe of the litigation. The transposition of these lessons
For discovery protocols, it should therefore be presented as analytical inference, and not as...
Command extracted directly from Mata v. Avianca.

3.5 Sanctions and the "intent to deprive" standard of Rule 37(e)

Rule 37(e) of the FRCP, as amended in 2015, establishes a specific regime.
sanctions for the loss of ESI that should have been preserved due to reasonable litigation
predictable. The rule distinguishes two levels of sanction: when the loss results from a failure to adopt measures.
reasonable grounds for preservation, but without proof of intent to harm the opposing party, the court may
to determine only curative measures proportionate to the demonstrated damage; when, however, it remains
demonstrated that the party acted with the specific intention of depriving the opposing party of the use of
In the absence of information in court, the court may presume that the missing information was unfavorable to the party.
spoliadora, instruct the jury in this regard or even dismiss the action or declare default (BOSTON
BAR ASSOCIATION, 2022; LOGIKCULL, 2018).
This distinction becomes relevant in the context of Shadow AI. Generative AI tools in use.
Customers may operate with retention policies defined by the supplier, not the organization.
user; therefore, the loss of prompts and outputs may result from third-party settings or practices.
The application of Rule 37(e) requires a concrete analysis, among other elements, of the existence of the duty to
preservation, control over information, prevention of actual ESI loss, and adoption of measures
reasonable. Thus, mapping relevant tools and communicating preservation obligations to
Employees can be part of a reasonable governance program, but the absence of such measures does not.
It must be described, in an automatic and abstract way, as sufficient grounds for sanctions.

4 METHODOLOGY

This article adopts a qualitative approach, of an exploratory nature and of a legal-doctrinal character and
documentary. The research is based on the analysis of three categories of primary and secondary sources:
(i) primary normative texts, notably the Federal Rules of Civil Procedure, especially the
Rules 26, 34 and 37, consulted in their current official version available from Legal Information.
Cornell Law School Institute; (ii) judicial precedents selected for their relevance.
paradigmatic for the evolution of technological e-discovery, notably Da Silva Moore v. Publicis
Groupe & MSL Group (2012) and Mata v. Avianca, Inc. (2023), obtained from public repositories.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

of judicial decisions; and (iii) specialized institutional and doctrinal literature, with emphasis on the
Publications from the Sedona Conference, a leading organization in the formation of professional consensus.
on e-discovery in the United States, supplemented by technical reports from the security industry.
information regarding the empirical extent of the Shadow AI phenomenon.
The approach is predominantly deductive, starting from the general normative structure of
proportionality standard, and then examine its application to a technological phenomenon.
specific — enterprise generative AI — and identify, through comparative analysis between the
The TAR regime and the emerging regime of generative AI: the technical and legal elements that must be considered.
to develop defensible ESI protocols. The research did not involve the collection of primary empirical data.
nor were there any interviews or original surveys conducted with legal professionals, limiting itself to
critical systematization of previously published documentary and jurisprudential sources, which is compatible with
the legal-dogmatic nature of the problem investigated, with its own time and scope limitations.
from a scientific article.

5. ANALYSIS AND DISCUSSION

5.1 The structural tension between proportionality and algorithmic opacity

The analysis undertaken in the previous sections allows us to identify a structural tension between the pattern.
proportionality of Rule 26(b)(1) of the FRCP and the opacity that characterizes both the use
The corporate landscape of generative AI is intertwined with the technical workings of the language models themselves.
The principle of proportionality presupposes that the parties and the court can reasonably estimate...
precision, volume, location, and cost of access to the relevant ESI universe. When parceling
This universe is produced using uninventoried Shadow AI tools; this estimate makes-
if structurally compromised — not due to bad faith on the part of the parties, but due to a governance deficit that
It precedes the litigation itself.
This finding suggests that the debate on proportionality in generative e-discovery cannot continue.
to be addressed exclusively on a procedural level, at the time of the meeting provided for in Rule 26(f);
It must go back to an earlier stage, that of corporate data governance, in which the organization
have already mapped, at least in general terms, the effective generative AI tools
used by their employees. In this sense, procedural defensibility in e-discovery
Generative AI is increasingly dependent on AI governance practices that
Traditionally, these were treated as matters of compliance and information security, and not of...

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google
RCMOS – Multidisciplinary Scientific Journal The Knowledge. ISSN: 2675-9128. São Paulo-SP.

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

procedural law — highlighting the growing convergence between these two areas.

5.2 From statistical defensibility to sociotechnical defensibility

A comparison between the consolidated defensibility regime for the TAR and the still nascent regime. Applicable to generative AI, this reveals a paradigm shift in the very nature of defensibility. required. In Da Silva Moore, the defensibility of technology-assisted review was established, largely through objective statistical metrics — accuracy, recall, and stability of iterative sampling process (DA SILVA MOORE V. PUBLICIS GROUPE, 2012). It was, therefore, of a defensibility that is essentially statistical, verifiable relatively objective, determined by experts and by the court itself.

The defensibility of generative AI, in turn, cannot rely solely on statistical metrics. performance, since the notion of "correct answer" is less determined in open-ended generation tasks. The focus is more on language than on relevance classification. Defensibility in this context is closer to... This article proposes to call sociotechnical defensibility: it depends not only on technical performance of the model, but also organizational factors, such as who reviewed the output, Based on which protocol and with what degree of independent verification? Mata v. Avianca offers a useful analogy regarding responsibility for human verification, although it is not a precedent. e-discovery (MATA V. AVIANCA, 2023). This approach suggests that ESI protocols Involving generative AI, balance technical requirements, when available and proportionate, with Documenting the human review workflows.

5.3 The asymmetric risk of Shadow AI for parties with different degrees of maturity in governance

The research also highlights a risk of a distributive nature: organizations with greater maturity. in information technology governance tend to be better positioned to map and documenting corporate use of generative AI, while smaller organizations or those with Less developed compliance structures face a disproportionate risk of exposure to sanctions under Rule 37(e), simply for not having the resources to monitor the phenomenon. Shadow AI is equally effective. This asymmetry relates to one of the very factors of proportionality provided for in Rule 26(b)(1) — the resources of the parties —, suggesting that courts In the future, they will need to consider with greater nuance the degree of technological maturity of each litigant.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

when assessing whether reasonable preservation measures have been effectively adopted, avoiding a pattern
The only one that disproportionately penalizes organizations with less technical and budgetary capacity.

5.4 Synthesis: elements for a new balance

The convergence of the previous analyses allows us to propose a normative synthesis: proportionality.
The use of generative e-discovery should not be interpreted as permission to ignore the universe of
data produced by generative AI due to its technical complexity, nor as
basis for indiscriminate discovery requests regarding any AI tool.
potentially used by an organization. The right balance lies in the requirement of
Proactive governance — mapping, custody, versioning, and algorithmic chain of custody, such as
proposed in Section 3.3 — combined with a reasonableness standard sensitive to the degree of maturity.
technological aspects of each part, just as already occurs, by analogy, with the evaluation of accessibility of
data sources under FRCP Rule 26(b)(2)(B).

6. FINAL CONSIDERATIONS

This article analyzed how the proportionality standard enshrined in Rule 26(b)(1) of
Federal Rules of Civil Procedure are being reshaped by the rise of artificial intelligence.
Generative intelligence in the corporate environment, with an emphasis on the phenomenon of Shadow AI. The research allowed
to identify that the traditional architecture of e-discovery — built on the pillars of preservation,
of the production and judicial control of proportionality — presupposes a degree of visibility regarding
the corporate data sources that Shadow AI tends to erode, creating a gap between the standard.
of "reasonable search" and the technological reality of contemporary organizations.
The analysis of Da Silva Moore v. Publicis Groupe (2012) and Mata v. Avianca, Inc. (2023) allowed
to contrast two distinct contexts: the first constitutes a milestone in the judicial acceptance of the review.
assisted by technology in e-discovery; the second, although not dealing with electronic discovery, offers
An analogy about the risks of uncritical trust in AI-generated content. The transition from ANT (Anti-Terrorism Financing).
For generative uses, it is necessary to expand the defensibility parameters, incorporating both controls.
technical aspects as well as organizational protocols for human verification. As a normative proposal, the
The article argues that ESI protocols negotiated in the Rule 26(f) phase should take into account, in accordance with
relevance and proportionality in the specific case, mapping of AI tools, preservation
of available prompts and outputs, model versioning, and production flow traceability.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

6.1 Study Limitations

This study has limitations that must be explicitly acknowledged. Firstly,
This research is primarily doctrinal and documentary in nature, without data collection.
primary empirical findings from courts, law firms, or legal departments
corporate, which limits the possibility of generalization regarding the true extent of the phenomenon of
Shadow AI is already the subject of ongoing litigation. Secondly, the object of analysis is limited to the system.
Federal civil procedure in the United States, not including, by methodological choice, a comparison.
systematic with other legal systems, including the Brazilian one, whose Code of Civil Procedure
and whose General Law on the Protection of Personal Data (LGPD) presents partially analogous provisions,
but not identical to those examined here. Thirdly, the topic under investigation is at a stage...
of particularly accelerated normative and jurisprudential development — as evidenced by
The circumstance that the "Primer" of the Sedona Conference on generative AI in e-discovery itself
it remained, until the conclusion of this research, in the draft phase subject to public comment (THE
SEDONA CONFERENCE, 2026) —, so the conclusions presented here should be
interpreted as a diagnosis of a transitional moment, and not as a description of a state of
a right that has been definitively established.

6.2 Suggestions for future research

For future research, it is recommended to develop empirical studies that examine, from
in a systematic way, subsequent judicial decisions that directly apply the parameters
proportionality and sanctions for spoliation of interaction records with generative AI,
allowing for empirical verification of the hypotheses formulated here on deductive grounds. It is suggested,
Furthermore, conducting comparative law studies between the North American e-discovery system and
the electronic evidence production institutes provided for in the Brazilian legal system, notably to
In light of the Code of Civil Procedure and the General Law on the Protection of Personal Data, in order to assess in
To what extent do the AI governance elements proposed in this article—tool mapping,
Prompt custody, model versioning, and algorithmic chain of custody—these could be...
incorporated into evidence production protocols in other jurisdictions. Finally, it is recommended that
investigation, from an interdisciplinary perspective that combines law and computer science, of
Objective technical auditability metrics for generative AI systems applied to litigation.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

capable of conferring upon this socio-technical defensibility a degree of objectivity comparable to achieved, in its time, by the statistical metrics of precision and recall employed in the review. assisted by technology.

REFERENCES

BOSTON BAR ASSOCIATION. The Impact of Recent Revisions to Fed. R. Civ. P. 37(e) — Electronic Spoliation. Boston Bar Journal, Boston, 2022. Available at: https://bostonbar.org/journal/the-impact-of-recent-revisions-to-fed-r-civ-p-37e-electronic-spoliation/. Accessed on: August 20, 2026.

CBIZ. How Shadow AI Compromises Security and Raises Risk and What to Do About It. [S. l.], 2025. Available at: https://www.cbiz.com/insights/article/how-shadow-ai-compromises-security-and-raises-risk-and-what-to-do-about-it. Accessed on: August 20, 2026. CLOUD SECURITY ALLIANCE. AI Gone Wild: Why Shadow AI Is Your Worst Nightmare. [S. l.], 2025. Available From: https://cloudsecurityalliance.org/blog/2025/03/04/ai-gone-wild-why-shadow-ai-is-your-it-team-s-worst-nightmare. Accessed on: August 20, 2026.

CYBERHAVEN. Shadow AI: How Employees Are Leading the Charge in AI Adoption — and Putting Company Data Risk. 2025. at [S. l.], Available at in: https://www.cyberhaven.com/blog/shadow-ai-how-employees-are-leading-the-charge-in-ai-adoption-and-putting-company-data-at-risk. Accessed on: August 20, 2026.

DA SILVA MOORE v. Publicis Groupe & MSL Group. 287 FRD 182. States District Court United States, Southern District of New York. Rapporteur: Judge Andrew J. Peck. New York, February 24, 2012. Available at https://law.justia.com/cases/federal/district-courts/new-york/nysdce/ in: 1:2011cv01279/375665/175/. Accessed on: August 20, 2026.

DIGITALWARROOM. The Simplest Guide to Understand FRCP Rule 26. [S. l.], 2019. Available From: https://www.digitalwarroom.com/blog/frcp-rule-26. Accessed on: August 20, 2026.

EVERLAW. TAR Case Law: Eighteen Rulings on TAR and Predictive Coding You Need to Know. [Place of publication unknown], 2024. Available at: https://www.everlaw.com/blog/ai-and-law/tar-predictive-coding-case-law/. Accessed on: August 20, 2026.

FINNEGAN, HENDERSON, FARABOW, GARRETT & DUNNER LLP. Effect of "New" Proportionality Limits in Amended FRCP 26. [S. l.], 2016. Available at: https://www.finnegan.com/en/insights/articles/effect-of-new-proportionality-limits-in-amended-frcp-26.html. Accessed on: August 20, 2026.

GLOBAL RELAY. The Complete Guide to the FRCP 26 Rules. [Place of publication unknown], 2024. Available at: https://www.globalrelay.com/resources/the-compliance-hub/rules-and-regulations/frcp-26/. Access Published: August 20, 2026.

IBM. What Is Shadow AI? [S. l.], 2025. Available at: https://www.ibm.com/think/topics/shadow-ai. Accessed on: August 20, 2026.

ITBREW. Will Shadow AI Get Worse in 2026? [S. l.], 2025. Available at:

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Machine Translated by Google

Year VII, v.2 2026 | Submission: October 6, 2026 | Accepted: October 21, 2026 | Publication: November 5, 2026

https://www.itbrew.com/stories/2025/11/14/will-shadow-ai-get-worse-in-2026. Accessed on: August 20. 2026.

JATHEON. FRCP 26(f) Rule: Meet and Confer Conference Checklist. [S. l.], 2026. Available at: https://jatheon.com/blog/frcp-26f-rule-meet-and-confer-conference-checklist/. Accessed on: August 20. 2026.

LEGAL DIVE. Shadow IT Surge Threatens Corporate Data. [Place of publication unknown], 2024. Available at: https://www.legaldive.com/news/shadow-it-surge-threatens-corporate-data-report/716754/. Access Published: August 20, 2026.

LOGIKCULL. A Guide to Revised FRCP 37(e). [S. https://www.logikcull.com/ [l.], 2018. Available at: learning/sanctions/frcp-37e-recent-caselaw. Accessed on: August 20, 2026.

MATA v. Avianca, Inc. 678 F. Supp. 3d 443. United States District Court, Southern District of New York. Rapporteur: Judge P. Kevin Castel. New York, June 22, 2023. Available at: https://law.justia.com/cases/federal/district-courts/new-york/nysdce/1:2022cv01461/575368/54/. Accessed on: August 20, 2026.

NEURALTRUST. What Is Shadow AI?: Risks, Detection & Prevention Guide for 2026. [S. l.], 2026. Available at: https://neuraltrust.ai/blog/shadow-ai-risks-detection-prevention. Accessed on: August 20. 2026. SECURITYWEEK. The Shadow AI Surge: Study Finds 50% of Workers Use Unapproved AI Tools. [S. l.], 2025. Available at: https://www.securityweek.com/the-shadow-ai-surge-study-finds-50-of-workers-use-unapproved-ai-tools/. Accessed on: August 20, 2026.

THE SEDONA CONFERENCE. The Sedona Conference Working Group Series. [S. l.], [sd]. Available at: https://www.thesedonaconference.org/wgs. Accessed on: August 20, 2026.

THE SEDONA CONFERENCE. The Sedona Conference on AI and the Law, Part 1: AI and Civil Litigation. Reston, 2024. Available at in: https://www.thesedonaconference.org/AI_and_the_law_part_1. Accessed on: August 20, 2026.

THE SEDONA CONFERENCE. Navigating AI in the Judiciary: New Guidelines for Judges and Their Chambers. Sedona Conference Journal, v. 26, Feb. 2025. Available at: https://www.ediscoveryllc.com/sedona-conference-navigating-ai-in-the-judiciary/. Accessed on: 20 August 2026.

THE SEDONA CONFERENCE. Primer on Generative Artificial Intelligence in Discovery (Member (Comment Draft). Phoenix, 2026. Available at in: https://www.thesedonaconference.org/sites/default/files/2026-04/Primer%20on%20GenAI%20in%20Discovery%20-%20Member%20Comment%20Draft%2004-30-26.pdf. Accessed on: August 20, 2026.

UNITED STATES. Federal Rules of Civil Procedure. Rule 26: Duty to Disclose; General Provisions Governing Discovery. Including Advisory Committee Notes to the 2015 Amendment. Legal Information Institute, Cornell Law School. [l.], 2015. Available at: [S. https://www.law.cornell.edu/rules/frcp/rule_26. Accessed on: August 20, 2026.

This article is published in open access under the Creative Commons Attribution license, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.