



[Home](#) > [BLOG](#) >

FROM AI-GENERATED COMPLAINTS TO HALLUCINATED PRECEDENTS AND AI-GENERATED EVIDENCE: EMERGING ISSUES IN DOMAIN NAME DISPUTES

[←](#) [→](#)

From AI-Generated Complaints to Hallucinated Precedents and AI-Generated Evidence: Emerging Issues in Domain Name Disputes

ANKUR RAHEJA / JULY 21, 2026 / [BLOG](#) / [LEAVE A COMMENT](#)

0

SHARES



FROM AI-GENERATED COMPLAINTS TO HALLUCINATED PRECEDENTS AND AI-GENERATED EVIDENCE: EMERGING ISSUES IN DOMAIN NAME DISPUTES

Article by Mr. [Stefaan Meuwissen](#), Knowledge Lawyer I.P. at Hogan Lovells International LLP

1. Executive Summary

Artificial intelligence has become a routine feature of proceedings under the Uniform Domain Name Dispute Resolution Policy (UDRP). Both complainants and respondents now use AI to draft pleadings, identify legal arguments, generate or support evidence, build websites and business plans, and produce automated domain name valuations. Used carefully, these tools can make the UDRP more accessible and more efficient. Used carelessly, they introduce real risks around accuracy, evidential reliability, confidentiality and accountability.

A consistent pattern is emerging from recent WIPO decisions: panels do not treat AI use itself as the issue. What matters is the quality, accuracy and probative value of the resulting submission or evidence. AI-generated content does not excuse a party from responsibility for inaccurate citations, unsupported assertions or misleading website content, and a polished AI-generated pleading cannot make up for the absence of evidence needed to satisfy the UDRP elements.

For panelists, AI can be a useful administrative tool, but human judgment and personal responsibility must remain central to the decision. The WIPO's Note on the Use of Artificial Intelligence (AI) in WIPO ADR Proceedings, built around accuracy, accountability, confidentiality, transparency and the principle that decision-making cannot be delegated to a machine (1).

2. Introduction

Generative AI can produce a professional-looking UDRP complaint or response in seconds, complete with what appears to be relevant doctrine, structured legal arguments and case citations. That has an upside: it may lower the practical barriers to participation and result in more respondents filing a substantive response instead of defaulting, and AI-assisted responses could eventually become the norm rather than the exception. However, it also has a downside. The fluency and apparent authority of AI-generated text can mask serious weaknesses, including fabricated authorities, inaccurate quotations, irrelevant legal analysis and factual assertions with no contemporaneous evidence behind them.

AI is also changing the evidential record parties rely on. Search results, website content, business plans and domain name valuations can all now be AI-generated, and in some cases the AI-generated material is itself part of the conduct in dispute, for example a website built almost entirely by an automated tool that is said to demonstrate bad faith. For counsel and panelists, going forward, the practical question is rarely whether AI was used. It is whether the resulting submission or evidence is accurate, relevant, verifiable and sufficient to establish what the UDRP actually requires.

3. Use of AI in Domain Name Cases: By Complainants

Complainants may use AI to draft complaints, identify legal principles and locate potentially relevant UDRP decisions. Done responsibly, this can improve structure and efficiency. Done without verification, it can produce a complaint built on citations that are inaccurate, irrelevant or simply invented.

In *WIPO Case No. D2025-4081*, the complainant cited sections of the *WIPO Overview of WIPO Panel Views on Selected UDRP Questions, Third Edition* that did not support the propositions advanced, and improperly quoted documents it had itself submitted as evidence (3). The panel could not tell whether the errors were the product of AI hallucination, a lack of diligence by the complainant or its attorney, or an attempt to mislead the panel, and treated all three explanations as equally troubling (3).

The consequences of relying on fabricated authority can go further than simply losing an argument. In *WIPO Case No. D2025-4174*, the complainant cited several UDRP decisions that do not exist and others that exist but do not support the propositions for which they were cited (4). The panel referred to well-publicized instances of lawyers submitting AI-generated work riddled with fabricated case citations to courts, and said that AI use was “the most charitable interpretation” available for the complaint before it (4). The complaint was denied, and the panel made a finding of Reverse Domain Name Hijacking (4). For complainant’s counsel, the lesson is direct: a complaint built on invented authority does not just fail on the merits, it can expose the complainant to an RDNH finding of its own.

AI output is also starting to appear as evidence, not just argument. In *WIPO Case No. D2025-5203*, the complainant submitted the result of a ChatGPT query about the historical association of the term “payshap” (5). The panel treated the geographically targeted, time-bound search results, together with the specific AI prompt and response, as the kind of evidence one would expect from a complainant seeking to rebut a respondent’s claim that it was unaware of the mark or that the term was purely descriptive (5).

This cuts both ways in practice. AI output can have real evidential value where the methodology behind it is transparent and specific enough to test, but it should be treated as supporting material similar to search engine results, rather than a substitute for primary or independently verifiable evidence. Complainants relying on an AI query as evidence should preserve the complete prompt, the full output, the date and any geographic or other parameters used, in the same way they would document a Wayback Machine capture or a trademark register search, and should be ready to explain the methodology if it is challenged.

4. Use of AI in Domain Name Cases: By Respondents

For respondents, generative AI may significantly lower the cost of preparing a structured response and identifying relevant UDRP principles. That may mean more respondents engage with a complaint instead of defaulting, which is generally a healthy development for the process per se. However, apparent professional presentation of arguments is not the same as professional representation or substantive merit, and panels are increasingly alert to the difference.

In *WIPO Case No. D2024-4224*, the panel considered that the response had likely been prepared with AI assistance. It made clear that provenance does not determine outcome, stating that “it is the content of the submission that matters” (6). At the same time, the panel noted that apparent AI use can invite closer scrutiny where the response does not squarely address the complaint’s allegations or the applicable UDRP elements (6).

In *WIPO Case No. D2025-2916*, the panel put the point more candidly: a generative AI model can produce something that looks, on the surface, like a professionally supported pleading, while inventing details to fit the prompt it was given. The panel considered that a response suggested to have been generated by an LLM was unlikely to inspire confidence in its veracity unless the panel could see the prompt and was satisfied that the output had been reviewed for accuracy and truth before submission (7).

These decisions point to a set of practical tells that panels, and opposing counsel, are starting to notice in AI-drafted pleadings. Tell-tales include a neutral, formal, third-person tone in a response filed by a respondent who is not represented by counsel, since self-represented respondents drafting their own responses tend to write in the first person; formulaic reasoning and wording clearly copied from the WIPO Overview 3.1 while not really being appropriate to the facts at hand are another indication; the word “defendant” used in place of “respondent”; case citations that exist but do not actually support the point being made; citations to sources that would not normally appear in a UDRP proceeding, such as treaties; and broad, unsupported claims of legitimate interests or preparations to use a domain name. None of these features are generally conclusive on their own, but together they may indicate that AI has been used.

The evidence gap is clearest in cases involving claimed business preparations. In *WIPO Case No. D2025-4065*, the respondent produced two apparently self-generated business plans but nothing else: no evidence of technical development, financing, marketing activity, expenditure or any other indicator of a genuine commercial undertaking, despite the scale of the business the plans described. The panel considered that the business plans were likely generated largely by AI and, in the absence of corroborating material, was not persuaded that the respondent had a bona fide intention to use the domain names (8). The same reasoning would apply to any unsupported claim of business preparation with or without AI involved, but AI makes it cheap to produce a plausible-looking plan with nothing behind it, which is exactly why panels are demanding more. In this regard, it is worth noting that the new WIPO Overview 3.1 now contains the following in section 4.2: *“In cases where arguments and/or evidence presented to the panel is created with the assistance of AI tools, the credibility of such arguments/evidence may be impacted by the inclusion (e.g., in an annex) by the parties of specific prompts and their results”*.

Reliance on an AI tool does not, by itself, create rights or legitimate interests. In *WIPO Case No. D2023-3906*, the respondent said that ChatGPT had suggested the disputed domain name for a group trip. The panel was not persuaded, both because the

explanation strained credulity on its own terms and because relying on a chatbot suggestion does not vest any rights in the resulting domain name, particularly where the name carries a risk of implied affiliation with the complainant (9).

Automated valuation evidence fares no better without an explained methodology. In *WIPO Case No. D2025-3938*, the respondent pointed to a high valuation supposedly generated by both AI and human assessment, including reference to “hundreds of predictive data points.” The panel found that such valuations, without detailed evidence of how they were reached, “mean very little” (10). The underlying question, of how much explainability an AI tool needs to provide before its output can be trusted, is not yet resolved in UDRP decisions, but it is a useful lens for testing any AI-generated valuation or search result a party submits (2).

AI-generated website content raises separate issues under both the legitimate interests and bad faith analysis. In *WIPO Case No. D2025-0217*, the disputed website had been built using GoDaddy Airo, a tool the registrar makes available to generate a basic website automatically once a domain name is purchased, with the domain owner able to edit or reject the content before it goes live (11). The panel held that the respondent remained responsible for the website’s content even though it had been generated automatically, regardless of how much the respondent had actually done to shape it (11). *WIPO Case No. D2024-5347* reached the same conclusion on similar facts: a respondent cannot disclaim responsibility for website content simply because AI generated it (12).

Two further decisions show how this plays out when the AI-generated content is itself deceptive. In *WIPO Case No. D2025-2612*, the panel found that a thin, AI-generated website with fabricated contact details, a template privacy statement and AI-authored articles under invented bylines was not the website of a bona fide business, though the complaint was ultimately denied on other grounds, a reminder that a suspect AI-generated website does not by itself relieve a complainant of its burden on the remaining UDRP elements (13). In *WIPO Case No. D2025-0332*, an AI-generated website that misappropriated third-party copyrighted images and offered services in

direct competition with the complainant was likewise found not to be a bona fide use, with the panel again stressing that responsibility for displayed content rests with the respondent regardless of how that content was produced (14).

The common thread across these decisions is therefore responsibility, not authorship. AI does not launder a website, a valuation or a business plan into something credible. If anything, an AI-generated artifact with no independent corroboration should invite more scrutiny, not less, precisely because it is so cheap to produce.

5. Use of AI in Domain Name Cases: By Panelists

AI may assist panelists with administrative, organizational and drafting tasks, but its use raises different considerations for decision-makers, who exercise adjudicative responsibility and have access to confidential submissions, evidence, deliberations and draft decisions. The WIPO Center's Note on the Use of Artificial Intelligence (AI) in WIPO ADR Proceedings recognizes the potential efficiency gains while emphasizing that AI cannot replace independent professional judgment. Procedural and substantive decisions, case-related communications and materials issued in a panelist's name remain the panelist's personal responsibility, regardless of whether AI was used (1).

Confidentiality is perhaps the most immediate concern. The Note treats information entered into AI tools as potentially accessible to third parties unless clear and verifiable safeguards are in place. Panelists should therefore avoid uploading pleadings, evidence, draft decisions, deliberations, personal data or other non-public case material where a tool's data-retention, training or confidentiality arrangements are unclear. The Note also encourages appropriate transparency and consultation with participants regarding AI use and allows panelists to seek disclosure from parties and counsel or establish procedural guidance for its use during proceedings (1).

WIPO's approach reflects a broader trend in dispute resolution. Guidance issued in the United Kingdom, Australia and the international arbitration community similarly treats AI as a supporting tool rather than a substitute for human judgment (15). The EU Artificial Intelligence Act follows the same principle: use of AI in these

circumstances is considered high-risk and while AI may support adjudicative decision-making, the final decisions must remain human-driven (16). Together, these developments point towards a common standard, AI may improve efficiency, but responsibility, confidentiality and independent judgment remain with the decision-maker.

6. Practical Takeaways

AI is likely to become an increasingly routine feature of domain name dispute resolution. Its value lies principally in supporting and not replacing legal analysis, evidential assessment and human decision-making.

For complainants and their counsel, every AI-generated authority, quotation and factual proposition should be independently verified against the original source. Where AI output is submitted as evidence, the prompt, date, parameters and complete response should be preserved, and the output should be supported by primary or independently verifiable material where possible.

Respondents should recognize that a seemingly polished AI-generated response cannot remedy the absence of evidence. Claims concerning preparations to use a domain name, business activity or domain-name value should be supported by objective contemporaneous and independently verifiable documentation and not just AI-generated business plans or other justifications. Respondents also remain responsible for content displayed on websites associated with their domain names, including content produced by automated AI website-building tools.

For panelists, AI may offer efficiencies, but confidentiality, independence and personal responsibility impose clear limits. Case materials should not be entered into AI systems unless adequate safeguards have been established, and substantive analysis and final decision-making must remain human-driven.

The emerging direction of UDRP practice is therefore pragmatic rather than prohibitive: AI use is not inherently improper, nor is AI-generated material inherently unreliable. Its weight and consequences will depend on accuracy, transparency,

corroboration and responsible human oversight.

Sources:

1. [WIPO Arbitration and Mediation Center, Note on the Use of Artificial Intelligence \(AI\) in WIPO ADR Proceedings.](#)
2. [WIPO Case No. D2025-4081](#), disabilitypridenm.org (David H. Bernstein, Panelist).
3. [WIPO Case No. D2025-4174](#), upscaleavenues.com (Robert A. Badgley, Panelist).
4. [WIPO Case No. D2025-5203](#), payshap.com (Andrew D. S. Lothian, Panelist).
5. [WIPO Case No. D2024-4224](#), euroveritas.com (John Swinson, Panelist).
6. [WIPO Case No. D2025-2916](#), arcelgroup.com (Andrew D. S. Lothian, Panelist).
7. [WIPO Case No. D2025-4065](#), moelisai.com and moelisglobal.com (Deanna Wong Wai Man, Panelist).
8. [WIPO Case No. D2023-3906](#), americanairlinesgroup.com (William F. Hamilton, Panelist).
9. [WIPO Case No. D2025-3938](#), onarsource.com (Andrew D. S. Lothian, Panelist).
10. [WIPO Case No. D2025-0217](#), bdö.com (Nick J. Gardner, Panelist).
11. [WIPO Case No. D2024-5347](#), asuriontech.net (Gary Saposnik, Panelist).
12. [WIPO Case No. D2025-2612](#), startablog123.com (John Swinson, Panelist).
13. [WIPO Case No. D2025-0332](#), topinsightglobal.com (Philip V. Marano, Panelist).
14. [UK Courts and Tribunals Judiciary, Artificial Intelligence \(AI\) Guidance for Judicial Office Holders \(2025\); Federal Court of Australia, Use of Generative Artificial Intelligence \(GPN-AI\) General Practice Note \(2026\); American Arbitration Association, Artificial Intelligence in Arbitration: Is There Room for AI Arbitrators? \(2026\); Chartered Institute of Arbitrators, Guidelines on the Use of AI in Arbitration \(2025\).](#)
15. [Regulation \(EU\) 2024/1689 of the European Parliament and of the Council of 13 June 2024 \(Artificial Intelligence Act\), recital 61.](#)

0
SHARES



Leave a Reply

Your email address will not be published. Required fields are marked *

Comment*

Name*

Email*

Website

Submit

[When .AI Cuts Both Ways: Applying UDRP Perspectives on Targeting – vol. 6.29](#)

July 21, 2026

[From AI-Generated Complaints to Hallucinated Precedents and AI-Generated Evidence: Emerging Issues in Domain Name Disputes](#)

July 21, 2026

[Chronological Barrier Defeats French Real Estate Firm Over 25-Year-Old Surname Domain – vol. 6.28](#)

July 14, 2026

[Andriy Khvetkevych](#)

July 10, 2026

[Descriptive AI-Branded Mark on Supplemental Register Dooms Complaint – vol. 6.27](#)

Contact the ICA

Contact

Name *

First

Last

Email *

Subject

Comment or Message *



I'm not a robot

reCAPTCHA

Submit