Minerva26

AI, Work Product, and the Protective Order Problem:What Morgan v. V2X, Inc. Means for Every Litigator

INTRODUCTION

Welcome to the Case of the Week segment of the Meet and Confer podcast. I'm your host Kelly Twigger. Thanks for being here today.

Well we are back to the topic du jour – generative AI. A few weeks ago we covered *U.S. v. Heppner* and *Warner v. Gilbarco* — the first two federal rulings on whether materials generated using publicly available AI tools are protected by the attorney-client privilege or the work product doctrine. Many of you heard us discuss those decisions at the University of Florida eDiscovery Conference, where we covered them on the case law panel with United States District Judge Xavier Rodriguez and Maria Salacuse from the EEOC. If you weren't there, or you haven't listened to that episode yet on the Meet and Confer podcast, go back and start there, because this week's case is the direct next chapter in that story.

I want to flag something that District Judge Rodriguez said at UF, because it will be a throughline in today's discussion. His point was that viewed through the lens of the law as it exists today, both *Heppner* and *Gilbarco* are correctly decided. But he said there needs to be a rethinking of the law on privilege in the age of AI. By the time we are done today, you are going to understand exactly why he said that and know that his thinking is 100% correct.

On March 30, 2026, United States Magistrate Judge Maritza Dominguez Braswell of the District of Colorado issued a ruling in *Morgan v. V2X, Inc.* that is, in my view, the most consequential AI-in-litigation decision we have seen yet. Now, yes, it's only the third one, but it's really well laid out. One of the challenges we had with both *Heppner* and *Gilbarco* was that they were pretty thin on analysis. It takes the questions that *Heppner* and *Gilbarco* left open, answers several of them directly, and raises a new question about access to justice and the economics of AI in litigation that I think is going to shape how courts and practitioners approach this issue for years. This is a decision you need to read — so please share it — and one you need to be talking about in your firms and with your clients. I hope to see a lot of client emails going out specifically discussing the *Morgan* decision.

Before I get into the case, I want to say something about Magistrate Judge Braswell, because context matters here. She is one of the most informed jurists writing on AI and the law right now. Despite being new to the bench in the last few years, she's an active member of The Sedona Conference Working Groups 1 and 13, co-chair of the District of Colorado's AI Committee, co-founded the Judicial AI Consortium, and authors The AI

Brief, a newsletter on AI's implications for courts. This is not a judge who stumbled into these questions. She has been living them, and it shows in every page of this opinion.

As always, Magistrate Judge Braswell — like every judge that we talk about on Case of the Week — is limited to what is before the Court in terms of making a decision. And as we go through this analysis today, I think you're going to see there are a number of places where if there was something slightly different presented to the Court, it would have dramatically changed how the Court ruled here.

As always, we'll have a link to the decision in Minerva26. Every decision in our database is available publicly, so you can click on the link to view the full text of the case. I built that feature because I wanted people to read the case law. And it infuriates me that much of our publicly available tax dollar case law is hidden behind firewalls. So you can click on that link to view the full text of the case. If you are a user of our Minerva26 platform, and you want to follow the generative AI case law that's coming out, we track all of those discovery decisions in the United States, including generative AI, and you can use the Generative AI issue tag in the platform to be notified immediately of new decisions on the issue by creating a dynamic cite list.

Let's get into today's decision.

THE CASE: MORGAN V. V2X, INC.

Morgan v. V2X is an employment discrimination case. Plaintiff Archie Morgan is a pro se litigant — meaning he is representing himself — who describes himself in the text of the decision as a "qualified Black American manager." He claims he was subjected to a hostile work environment, terminated based on his race and national origin, and retaliated against for opposing sexual harassment and engaging in protected whistleblowing. The defendant, V2X, says he was fired for legitimate, non-discriminatory reasons following a thorough workplace investigation involving more than 30 witnesses.

Sound familiar? Well, it should. Just like Sohyon Warner in *Gilbarco*, Morgan is a pro se plaintiff in an employment discrimination case who has been using AI to help him litigate. And just like in *Gilbarco*, that AI use has become its own battlefield.

Before I get to the motion, some very important context here. The parties in this case had already entered into a stipulated protective order. It allowed either party to designate information as Confidential and implicates the privacy or business interests of the parties — things like medical information, private personnel records, trade secrets, and proprietary business information. Under that stipulated protective order, Confidential Information could not be disclosed except under specifically enumerated circumstances. Now here is something the opinion makes clear and I think deserves emphasis before we go any further: BOTH parties are using AI in connection with their litigation work. This is not a case where one side is accused of doing something unusual or outside the bounds. Both sides are using AI. They just disagree on how it should or should not be used in connection with Confidential Information and whether what tools they're using should be disclosed. Keep that symmetry in mind.

Now, I want to walk you through how this dispute actually came to a head because it's really good context and it also really shows you how one party seeks to leverage things over another in litigation. Oftentimes, I really feel like courts need to address these kinds of situations more. Doesn't seem to be addressed much here. The

procedural backstory here tells you something important about how corporate defendants are thinking about AI exposure — and how high the stakes are.

This motion was filed by V2X shortly after Morgan moved to compel the production of an insurance policy. That policy was four months overdue. Morgan's position was pretty direct: V2X was refusing to produce the insurance policy as a condition of getting him to agree to do two things — disclosure of which AI tool he was using, and amendments to the protective order restricting his AI use. He said V2X was holding a four-month overdue discovery obligation "hostage" to extract concessions on unrelated issues. And he framed it in terms that Magistrate Judge Braswell ultimately uses in her own opinion: V2X was creating an "unfair technological gap" by barring a pro se litigant from using modern analytical tools while V2X's counsel maintained its own proprietary AI and cloud-based systems. Important distinction here. V2X has its own systems that it's using, internal, likely not consumer grade. And the pro se litigant is trying to leverage AI for the benefit of himself as well.

The court struck Morgan's motion to compel for failure to follow discovery dispute procedures which are very specific in the District of Colorado, but set a discovery hearing to address his concerns. And then, the morning of that hearing — surprise, surprise, surprise — V2X's counsel emailed chambers to say that the parties had resolved the dispute — V2X had produced the insurance policy — and asked to vacate the hearing. The Court vacated the hearing.

V2X produced the document. The hearing went away. But V2X pressed forward with the motion to amend the protective order and compel disclosure of Morgan's AI tool. The AI restrictions and disclosure question was serious enough to V2X that it had apparently been willing to condition a routine discovery production on getting those concessions. When you understand that, the rest of this opinion really lands a lot differently.

One more thing about this motion before we get to the analysis. The opinion does not cite any legal basis for V2X's motion to amend the protective order. There is no standard articulated, no threshold showing V2X is required to meet, no rule cited that authorized it to reopen and modify a stipulated protective order that was already in place, and no factual basis that's laid out in the decision as to why V2X was coming to this. That is a gap in the record. And I think the reason that the gap exists is this: Morgan did not oppose the motion to amend. He opposed V2X's specific proposed language and offered his own competing language instead. The Court even specifically states that both parties agree to the fact that the protective order should be amended. But he did not challenge whether V2X had any right to seek the amendment in the first place. A court cannot rule on arguments that a party does not make. So if Morgan had challenged the basis for the motion itself — if he had said, show me the standard, show me what gives you the right to reopen the stipulated order — Issue Two of this case might have gone very differently. Morgan is a pro-se litigant and he was using AI.

Now, if you use AI, you know that if you put in a motion for something which is a publicly available document, so not confidential information, the AI will respond to that motion. It will NOT, unless you specifically ask it to, tell you what your options are in litigation other than responding to it — for example, opposing the motion here. That's something to keep in mind, because AI is like any other technology but harder – you have to be using it extensively to understand the ins and outs of it and what your potential legal arguments are around it within the legal framework that we have currently.

Magistrate Judge Braswell takes on both issues presented here: work product protection and the protective order language. And the analysis she gives us is the most thorough judicial treatment of AI and privilege in a civil case that we have so far.

ISSUE ONE: WORK PRODUCT PROTECTION AND A PRO SE LITIGANT'S USE OF AI

Let's start with the work product question. This is a critical issue for litigators in this case, and I want to walk through it carefully because the analysis has direct implications for how you advise clients who are using AI right now.

Magistrate Judge Braswell frames her analysis around two specific questions, and I want to quote her framing directly because it tells you exactly what is at stake: "This dispute raises two such questions: (1) to what extent will work product protections apply to a *pro se* litigant's use of AI, and (2) to what extent should a protective order expressly restrict the use of AI? The Court addresses each question in turn." Two clean questions. Now, this is a judge who knows how to write, and those questions map precisely onto the two issues practitioners are facing in their own cases right now, especially after *Heppner* and *Gilbarco*.

First, a quick grounding on the doctrine around the work product protection. For those who are not familiar, the work product protection is codified in Federal Rule of Civil Procedure 26(b)(3). It protects documents and tangible things prepared in anticipation of litigation or for trial by or by a party for a party or its representative. The doctrine has two tiers. The first tier provides qualified protection for ordinary work product — factual materials, research, documents assembled for the case.

That protection can be overcome if the opposing party shows substantial need and an inability to obtain the equivalent without undue hardship. The second tier provides near-absolute protection for opinion work product — the mental impressions, conclusions, opinions, and legal theories of a party or counsel. That heightened protection is essentially impossible to overcome. Rule 26(b)(3) traces its origins to the Supreme Court's 1947 decision in *Hickman v. Taylor*, which Magistrate Judge Braswell cites, and was specifically broadened in the 1970 amendments to the Federal Rules of Civil Procedure to extend protection beyond attorneys to parties themselves. That history really matters here.

The threshold question that Magistrate Judge Braswell has to answer is: does Rule 26(b)(3) apply to a pro se litigant's use of AI at all? And the Court notes that the Tenth Circuit, which is precedent here because Colorado sits in the Tenth Circuit, hasn't weighed in on this. And it's not a trivial question at all. Work product protection in the classic sense is about protecting the lawyer's thought processes — their theory of the case, their selection of what matters, their preparation for trial. So when a pro se litigant uses AI, there is no lawyer in the room. Does the doctrine still apply? We address this in the *Gilbarco* case.

Magistrate Judge Braswell says yes, and she works through it systematically. The plain text of Rule 26(b)(3) broadly protects materials prepared "by or for another party" — not merely by counsel. The Rule's history reinforces that: the 1970 amendments were specifically designed to extend protection beyond attorneys' materials to materials prepared by or for a party. Courts across the country have applied the rule to pro se litigants for decades, and there is genuine consensus on that point. So we've got *Gilbarco* and now we've got *Morgan* saying the same thing.

But then she says something I think is key and that goes beyond what the earlier decisions said. Magistrate Judge Braswell writes that the importance of extending these protections to pro se litigants is "magnified in the context of AI—one of the most powerful knowledge tools ever to become available to the masses." The reason she gives is that a pro se litigant has to be both a party and advocate at the same time. That is an enormous disadvantage in litigation against a represented corporate defendant. And for the first time in history, AI may actually make that dual role manageable. I mean, think about that. If you're the party, you are caught up in the emotion of what's happening. The allegations that you are making feel very personal to you to be able to separate that and think objectively about legal strategy going forward and have to do both at the same time. That's an extremely difficult position for a pro se litigant to be in. But a lot of pro se litigants have claims that they can't afford to bring because they cannot afford to hire a lawyer. So their choice is pro se or nothing.

Magistrate Judge Braswell found that conditioning work product protection on the involvement of counsel finds no support in the text of Rule 26 and would further disadvantage unrepresented litigants — litigants who are already held to the same substantive standards as represented parties. If a pro se litigant cannot protect the thought process they develop with AI in the same way an attorney can protect their own work product, the doctrine is doing the opposite of what it was designed to do.

So work product applies here.

The next question Magistrate Judge Braswell addresses is: does using an AI platform — which is technically a third-party system that collects, stores, and in many cases uses your data to train — waive that protection? This is V2X's core argument, and it's the same argument that the government made in *Heppner.* The logic is that you voluntarily disclosed your litigation preparation to a third party. Voluntary disclosure to a third party traditionally waives privilege. Therefore, work product protection is gone.

Magistrate Judge Braswell really digs into that argument and takes it seriously. And her answer starts with whether using an AI platform creates the kind of privacy-defeating disclosure that should matter under the work product doctrine.

She looks at two cases from the Fourth Amendment world — the Supreme Court's 2018 decision in *Carpenter v. United States* and the Sixth Circuit's decision in *United States v. Warshak* — not because they govern here, but for the principle they articulate: the mere fact that information passes through a third-party system does not automatically extinguish all privacy expectations. In *Warshak*, the Court held that email subscribers have a reasonable expectation of privacy in the contents of their emails even though a commercial internet service provider holds them. The intermediary's access alone does not eliminate privacy. Interestingly, now that Google will tell you on its free accounts that it can use your emails to train its servers, would that change the analysis? *Carpenter*, under Magistrate Judge Braswell's analysis, extended that reasoning to cell-site location data: the fact that a third party holds the information does not mean the person who shared it forfeited all privacy interests in it.

Those two decisions make great sense. The question is now, how do we extend them?

Then Magistrate Judge Braswell asks a question that every litigator needs to focus on. It is rhetorical, but it is the right question, and it is the logical extension of V2X's argument. She writes:

> *Today, nearly all electronic interaction passes through third-party systems. Google, for example, hosts millions of accounts, and by extension, has access to millions of messages, emails,*

> *documents, videos, and more. Moreover, we now know that our phones, computers, in-home smart devices, and other electronics, collect information about us to offer more bespoke services. Does that mean that anyone with a Gmail account has forfeited all rights to confidentiality and privacy?*

If you accept V2X's premise — that disclosure to an AI platform destroys confidentiality and waives protection because a third party has access — you have to accept that premise about Gmail too. About your phone. About your smart speaker. About every piece of modern technology that processes your data. I've got a smart thermostat on the wall here in my office. Is that subject still to privacy? That is not a sustainable legal rule. And that is exactly the point she is making.

She then makes a point that I think is particularly significant and that is supported by actual research, not just intuition. She says the case for privacy is STRONGER with modern AI tools than with email or traditional cloud storage.

AI platforms are not passive repositories. They are specifically designed to engage. They simulate empathy. They invite candid disclosure. They interact in a way that feels intimate and genuine, scarily so sometimes. And social science research confirms that people share more personal and sensitive information with AI chatbots than they do with any other digital tool, often without appreciating what happens to that information once shared. Here is what Magistrate Judge Braswell writes on the subject:

> *Unlike a general-purpose search engine, which passively returns results, many modern AI platforms are specifically designed and trained to engage. They invite candid and significant disclosure of information, including sensitive information. They simulate empathy, foster trust, and interact in a way that feels genuine and intimate.*

The implication of that language is that an AI user's reasonable expectation of privacy in that interaction is higher, not lower, than in a traditional digital communication. The platform is designed to make you feel safe sharing. That design has legal consequences.

V2X also argued that Morgan waived work product by disclosing information to the AI platform, and under established Circuit authority, work product protection is waived by disclosure to an adversary or in circumstances that substantially increase the likelihood that an adversary will obtain the materials. That is a targeted standard — not triggered by any third-party disclosure, but specifically by adversarial exposure. And like United States Magistrate Judge Patti in *Gilbarco*, Magistrate Judge Braswell found that AI platforms are not adversaries. Uploading litigation analysis to Claude or ChatGPT does not put it in V2X's hands. Absent a legal process compelling the AI provider to produce the data, it does not go anywhere. And more on that later, because the form in which that provider could actually give it back to you is somewhat suspect here. Magistrate Judge Braswell cites Magistrate Judge Patti's reasoning in *Gilbarco* specifically on this issue: "AI programs are tools, not persons, and the disclosure is therefore not to an adversary or in a way likely to get in an adversary's hands."

Now, as I mentioned, here is where I want to flag something that every litigator in this audience needs to hear. Magistrate Judge Braswell distinguishes *Heppner* on two grounds, and the second one is CRITICAL. The first is clean: *Heppner* was a criminal case; this is civil, governed by Rule 26(b)(3)'s broad text. District Judge Rodriguez made that distinction at UF.

But the second distinction is the one that should keep you up at night as outside counsel. In *Heppner*, there was a structural gap between the represented defendant and his attorneys because Heppner used Claude, an AI platform, entirely on his own initiative, without any direction or involvement from counsel. Magistrate Judge Braswell is not disagreeing with that result. She's saying it doesn't apply here because Morgan IS his own counsel. In the pro se context, there is no gap. But the implication runs in the other direction too: if a REPRESENTED party uses AI independently of their lawyer — not at counsel's direction, not as counsel's agent — the materials generated may not reflect counsel's strategy and may not merit work product protection. That is what *Heppner* held, and *Morgan* does not disturb that holding.

So think about what that means practically. A corporate executive under investigation who starts using AI on their own to develop their defense narrative, without their lawyers directing that work, may be in exactly the same position as Heppner. No privilege. No work product. If you represent a client you think might be doing that, this is a conversation you need to have now, before the materials are seized or subpoenaed. White collar counsel should be aware of that and instructing their clients as such.

All right, back to our facts here in *Morgan* and the final word that work product applies to Morgan's AI use, and it is not automatically waived because he IS his counsel. The next question is how far the protection reaches.

Morgan pushed for maximum scope. He did not just want to protect the prompts he wrote, the analysis he generated, or the interactions himself. He wanted to protect the NAME of the AI tool he was using — arguing that tool selection itself is a mental impression that reveals strategy, analytical approach, and litigation resource allocation.

And here's the specific language Magistrate Judge Braswell uses to explain why he lost on that: "you have not demonstrated that identifying the tool itself will reveal your mental impressions or legal strategy, and Defendant needs that information to assess whether Confidential Information was compromised." Now read those two clauses carefully, because they're two different things. The first is about Morgan's failure to carry his burden in showing that the name of the tool carries any sort of strategic product implications. The second is about V2X's stated rationale for wanting the information. And I'll come back to that second clause in a minute, because I think it deserves scrutiny.

On the first clause: Morgan lost because he offered only conclusory assertions. If you are a regular listener to Case of the Week, you know this is a theme I return to constantly. Conclusory allegations are never going to be enough on these discovery issues. You have to do the work to create a factual and legal record of why you should get what you want.

But here's what I want to push on, because I think there's a much stronger argument here than Morgan made, and it is one that litigators need to have in their toolkit. Nobody uploads documents to an AI platform passively unless you're just asking for summarization. The act of uploading is almost always accompanied by direction. You are telling the system what to DO with the documents. Even when I want something summarized, I know what I want the specific pieces of that summary to look like, whether it's structure or whether it's content. You're telling the system what to do with the documents when you upload with direction. "Analyze this personnel record for evidence of pretextual treatment" is an example. "Compare these two witness statements for internal inconsistencies." That's another statement. "What questions does this timeline raise about the sequence of events?" Those prompts ARE your mental impressions. They reveal your case theory, what you think matters in the evidence, where you believe the weaknesses are. They are arguably the purest form of

opinion work product that exists — your legal thinking articulated in real time, directed at the specific facts of your case, some of which may come in the form of Confidential Information. Here, it may have been possible that V2X was concerned that Morgan would upload the very insurance policy it refused to provide for 4 months and ask it questions that would help his case. Now, I don't know that and it's not in the decision, I'm just surmising. I also wonder why an insurance policy would be confidential though, too. So we'll get to that issue in a little bit.

The tool name that Morgan sought to protect is one step removed from the prompts. But it is not disconnected. Different AI tools have different capabilities and different analytical strengths. If you use more than one AI tool, you know exactly what those are. We can spend another hour talking about those separately. Knowing which tool you used helps you. It tells you something about the analytical framework that is applied. A sophisticated opponent could use the same tool with similar prompts and begin approximating what the other side found. The argument Morgan could have made connects those dots: I selected this tool because of specific capabilities that reflect my analytical approach to specific evidence in this case. Maybe I've trained it with specific skills or other background knowledge that is not confidential to provide better evaluation. Knowing which tool I used is the first step toward understanding my case theory. Those are potential arguments that Morgan might have used, although I concede that we could spend a good part of the day going back and forth on how viable those arguments are and what other arguments exist. I'd love to hear them from you, so please drop them in the comments.

The way to make that argument well is with specifics. A sworn declaration explaining not that "tool selection reveals strategy" — which is exactly the conclusory framing the Court rejected — but HOW it reveals strategy: what capability this tool has that drove the selection, what aspect of the evidence that capability was directed at, and why knowing the tool name would allow V2X to approximate the analysis. An analogy to *Sporck v. Peil*, the Third Circuit case where counsel's grouping of documents out of thousands produced in discovery was protected because the selection reflected mental impressions. That's one argument that could have been made. A request for in camera review of the tool's specific features before ruling on disclosure. Morgan didn't do any of that. He declared it would reveal strategy and moved on. And that's not going to be how you carry a work product burden. Again, Morgan is a pro se litigant here. So we're asking a lot from a pro se litigant that, quite frankly, many litigators wouldn't know to make an argument about.

Now, the second clause in the Court's language: "Defendant needs that information to assess whether Confidential Information was compromised." This is framing the disclosure requirement as a security question, not a strategic one. And I want to push back on that framing directly, because I'm not sure that it holds up.

If V2X's concern is genuinely about security — was its Confidential Information uploaded to a platform that could expose it — then the tool name alone does not answer that question. Knowing Morgan used Claude tells nothing about whether he was using the free consumer tier, a paid individual account, or an enterprise account with data processing protections. I went back and looked at the language of the order, and it only orders disclosure of the name of the tool. It doesn't order disclosure of any plan type. The security risk is entirely a function of the account type and the applicable terms of service, not the platform name. A more targeted order would have required Morgan to produce documentation of the account settings and applicable terms — that is what actually answers V2X's stated concern here.

We have also never required parties to disclose the software tools that they use to manage litigation data before. Confidential information, in fact. Document review platforms process your documents. Case

management software processes your communications. Microsoft 365 processes everything. None of that has ever required disclosure. The principle that using a digital tool with a third-party provider creates a disclosure obligation is a new legal requirement, and it has implications that go well beyond Morgan's choice of an AI tool. That's where I have some issues with Magistrate Judge Braswell's decision here.

ISSUE TWO: CRAFTING AI-SPECIFIC PROTECTIVE ORDER LANGUAGE

The second issue in *Morgan* is one that *Heppner* and *Gilbarco* never had to address: what should a protective order say about AI? This is new ground. No court had provided this level of guidance before. Magistrate Judge Braswell analyzes both parties' proposed language in detail here, rejects both of them, and writes her own. What she produces is the most comprehensive judicial framework for AI use in discovery to date in a protective order. And I want to get into the details, because the implications are broader than they might first appear.

V2X proposed language that named specific platforms — such as ChatGPT, Harvey. Magistrate Judge Braswell was direct about the problem with that: the language was over-engineered, clearly drafted around V2X's own existing enterprise contracts, not around the needs of the existing litigation or the practical reality of Morgan's situation. She was also very sharp in criticizing V2X about the fact that its proposed provision still referenced "Google's Bard" — a platform that has been rebranded as Gemini more than a year ago — and that an error like that in a proposed order tells you something about the care with which it is drafted.

Morgan's proposal permitted AI use as long as the platform operated in a "secure, closed-circuit environment" and didn't use data for LLM training. Magistrate Judge Braswell found that too narrow and largely off-target. Now that's key. She found that to be too narrow. So V2X was too broad, and plaintiff's proposal was too narrow. A secure closed-circuit environment addresses unauthorized access — cybersecurity risks from bad actors. But that's not the primary concern with consumer AI platforms. The concern is what the platform itself does with your data in the normal course of operations: storing it, using it to improve the model, disclosing it to third parties. Morgan's proposal mostly missed those risks.

So she wrote her own language. And here's the language that Magistrate Judge Brasswell laid out:

> *No party or authorized recipient may input, upload, or submit CONFIDENTIAL Information into any modern artificial intelligence platform, including any generative, analytical, or large language model-based tool ("AI"), unless the AI provider is contractually prohibited from: (1) storing or using inputs to train or improve its model; and (2) disclosing inputs to any third party except where such disclosure is essential to facilitating delivery of the service. Where disclosure to a third party is essential to service delivery, any such third party shall be bound by obligations no less protective than those required by this Order. In addition, the AI provider must contractually afford the party or authorized recipient the ability to remove or delete all CONFIDENTIAL information upon request. A party intending to use AI that it contends meets these requirements must retain written documentation of these contractual protections.*

Now let's break this down element by element, because each piece is targeting a specific risk. And whoa, but this language has some implications.

First, the prohibition on training use is the most critical element. Every major consumer AI platform — the free and low-cost versions of ChatGPT, Claude, Gemini, Perplexity — they all reserve the right to use your inputs to train and improve the model. That is how these platforms are built and paid for.

The prohibition on third-party disclosure with a pass-through obligation addresses the subcontracting reality of modern AI infrastructure. These platforms don't operate in isolation. They use cloud providers, sub-processors, and infrastructure partners. Your data may touch multiple third-party systems before a response comes back to you. Magistrate Judge Braswell's language requires that any such third party be bound by obligations at least as protective as the order itself. Now that's a meaningful requirement that most consumer providers cannot satisfy. But it's also going to place a huge obligation on anyone using tools like this under a protective order with this language in it. We'll talk about that a little bit.

Next, the right to delete is the third leg. If Confidential Information has been uploaded to a system, the party must be able to get it out — to remove or delete it upon request. The written documentation closes the loop: you cannot assert that your tool meets the standard, you have to have paper to prove it. Ironic that we're using an AI tool and you require paper.

That second hurdle to me is really problematic and we're going to talk about that.

Now, Magistrate Judge Braswell is completely candid about what this standard means in practice. And I want to read to you what she says, because she says it directly and it matters:

> *The Court recognizes that practically speaking, and in light of the current state of AI, this provision will (at least for now) bar the parties from using most, if not all, mainstream low-to-no-cost AI to process Confidential Information. This type of restriction disadvantages pro se litigants. Enterprise-tier AI accounts that satisfy these requirements may be available only through organizational procurement processes, or at costs that a pro se litigant is unlikely to bear. But the Court cannot ignore the real risks associated with mainstream tools that persistently collect and store data and could compromise confidentiality.*

She also then clarifies the scope:

> *The Court does not intend to leave pro se Plaintiff without the benefits of AI. Modern AI tools may be used in many ways that do not involve uploading Confidential Information, and nothing in this particular Order restricts those uses. What this Order requires is that Confidential Information not be entrusted to platforms that lack the contractual safeguards described above, regardless of the sophistication or apparent trustworthiness of the tool.*

Does Magistrate Judge Braswell's language mean that we have to disclose every platform we use to upload documents for review or analysis, along with its security requirements? Because that is what this standard seems to suggest in part two. If you are using a document review platform that has AI features — and nearly all of them do now — and you are uploading Confidential Information to it, does that platform need the contractual protections she describes? Under a strict reading of her language, the answer may be yes, and under the second part of her language, it means that every single third party down the line for every single tool has to meet the standards. Who is going to do that analysis on every tool?

Many cloud platforms are hosted in Google Cloud or AWS or Azure. Are we going to get all of that language together? Is that going to be provided consistently? Are we going to preclude litigation teams from using new tools until they go through the workflow of having this security documentation in place so they can comply with this language of a protective order? Or is it going to essentially mean that parties don't want language regarding AI in their protective orders because it subjects them to all kinds of burdens and potential productions that they don't want to make?

Now that is a fundamentally different world than the one litigators have been operating in. We have never had to audit and disclose our document review infrastructure as part of discovery practice.And there's nothing in the Federal Rules of Civil Procedure that provides that we should. We've had arguments even over whether or not Rule 34 requires that we tell the other side we're using technology assisted review or AI-assisted review, right? There's nothing in that rule that requires that. There's been many arguments in case law about that. If this standard extends to all litigation technology that has AI baked in, the compliance burden becomes enormous and falls hardest on the parties who can least afford to meet it. But it's also going to have a significant impact on any complex organization that has multiple AI tools.

Now I want to raise two additional problems with this language that the opinion does not address, and that are going to generate significant litigation.

The first is a technical problem. When you input data into an AI platform, that data is processed through what are called vector embeddings — mathematical representations of your content that are stored in a vector database. The data does not persist as the original document in the way a file stored on a server does. It is transformed. It is embedded in a mathematical structure that the model uses to generate responses. The “right to delete” that Magistrate Judge Braswell builds into her standard assumes that the data can be cleanly extracted and removed. But for data that has been processed through a large language model’s infrastructure, that technical deletion may not be complete or verifiable in the way a court order requires. Enterprise accounts with dedicated data processing agreements typically do have genuine deletion capabilities with audit trails. Consumer accounts may offer a delete button, but whether that satisfies a court-ordered protective obligation is an open question that Magistrate Judge Braswell does not address. I think by limiting the ability to put confidential information in platforms that don't have these security requirements, she feels that she is addressing that completely. But what about for non-confidential information? Is that not a problem anymore?

The second problem is the breadth of the language itself, and I think this has the most far-reaching implications. The order covers “any modern artificial intelligence platform, including any generative, analytical, or large language model-based tool.” Read that literally. Westlaw’s AI-assisted research uses large language model technology. Relativity’s AI review features process your documents analytically. Microsoft 365 Copilot is a generative AI tool baked into Word, Outlook, and Teams. Grammar tools. Smart drafting features. Email summarization. AI is not a separate product category anymore. It is baked into the infrastructure of all of our modern legal practice. Every tool that we already use, all of our cloud-based tools are building AI features in and providing them at the contract prices we're already using. If you have a Gmail account, you've been given AI features.

If you read Magistrate Judge Braswell’s standard strictly, all of those tools require the same contractual protections: explicit prohibition on training use, third-party flow-down obligations, and verified deletion rights. And the only parties who can satisfy that standard across the board are the ones with enterprise contracts, institutional procurement processes, and IT infrastructure capable of reviewing and managing those

agreements. Large firms with enterprise Microsoft 365 agreements and purpose-built legal AI deployments — yes, probably. Solo practitioners, small firms, public defenders, legal aid organizations, pro se litigants — no.

Magistrate Judge Braswell identifies the access-to-justice problem in her footnote, and I want to give her credit for putting that in the record. She writes:

> *This highlights a growing problem in the age of AI: as large firms pour thousands of dollars into enterprise-grade AI and make their use of AI more secure, efficient, effective, and powerful, how will a pro se litigant or a litigant who cannot afford big-ticket legal services and better AI keep up?*

But the problem her own language creates is not limited to pro se litigants. It's a structural advantage for any party with enterprise infrastructure over any party without it. She is not drawing a line between people who use AI and people who don't. She is drawing a line between people who can afford enterprise software and everyone else. That is a much larger and more systemic problem than a single footnote can carry. Any small business provider could use customer, consumer-based AI, but not have enterprise grade AI. Any mom and pop shop in a manufacturing industry may not use AI at all, but a defendant, you know, in a supply chain argument may be able to leverage enterprise AI to do all kinds of things in support of their defense that the mom and pop shop can't do. So we're creating another barrier here.

She does clarify that the order does NOT prohibit all AI use. Parties can use AI for research that has its own set of problems, drafting, and analysis of non-confidential materials. The restriction is specifically on feeding Confidential Information into the platforms that lack the contractual protections. That is an important distinction. But the practical effect — for a pro se litigant in an employment case who has no enterprise account and whose case involves confidential employment records — is that AI is effectively off-limits for the materials that matter most. Now, there's so many discussions you could get in here on the confidentiality issues, but that's for another day.

She also tells the parties directly to avoid over-designating documents as Confidential, in part because of the burden this creates. That caution is very meaningful and I will come back to it in the takeaways.

HOW THE LAW IS EVOLVING: HEPPNER → GILBARCO → MORGAN

Let's talk for a minute about how the law is evolving.

I want to zoom out and look at these three cases together — *Morgan*, *Gilbarco*, and *Heppner*, because what's happening here is important and it moves fast. And I know our Case of the Week segment is going long this week, but this is really important stuff that you're going to want to pay attention to.

We have three federal decisions issued over a span of seven weeks, from three different courts, all grappling with the same fundamental question: what legal protections apply when a litigant uses generative AI in connection with their case? And when you read them in sequence, you can see the law starting to find its shape.

Heppner was the most restrictive outcome, and it got the most attention. But look carefully at the facts: a REPRESENTED criminal defendant, charged on $150 million in securities fraud, using Claude ON HIS OWN without any direction from counsel, after he already knew he was the target of a federal investigation, on a platform whose privacy policy explicitly flagged potential government disclosure, with the materials then seized

pursuant to a search warrant by the government. Every single factor cut against protection. United States District Judge Rakoff applied existing law to those facts and reached the right result. But he was also writing in a very narrow lane.

Gilbarco arrived the same day as Heppner's bench ruling but got far less attention. PRO SE plaintiff, civil case, materials created during the discovery period, defendant seeking to compel production of her AI interactions without any evidence she had violated the protective order. Magistrate Judge Patti said no, the materials are work product, AI tools are TOOLS not persons, disclosure to a tool is not disclosure to an adversary, and no court has endorsed the theory that AI waives work product in modern drafting environments. His closing line will be quoted for years:

> *"In the end, both sides of this dispute seek to obtain each other's thought processes, while shielding their opponent from discovery of their own. The Court will uphold the protections afforded the thought processes and litigation strategies of both sides and will order production of neither."*

Morgan is where the two lines come together and the analysis deepens. Magistrate Judge Braswell explicitly distinguishes *Heppner* — she doesn't ignore it or wish it away. She says: criminal versus civil, and represented party acting alone versus a pro se litigant who IS simultaneously the party and the advocate. Those are real distinctions, and they are the right ones.

She aligns with *Gilbarco* on work product and waiver, and then adds the privacy analysis from *Carpenter* and *Warshak* to give the reasoning more foundation. And then she goes further than either prior case in two important ways.

FIRST: she draws the internal limits of the protection. You can't just invoke work product and stop there. You have to show HOW the specific disclosure you're resisting — in this case, naming the tool — would actually reveal your mental impressions or strategy. That is a meaningful constraint that practitioners need to internalize.

SECOND: she addressed the protective order question, which neither Heppner nor Gilbarco had to answer. She gives us a standard — a contractual prohibition on training, prohibition on third-party disclosure, right to delete, and written documentation — and she tells us honestly that consumer AI almost certainly doesn't meet that standard. That is the most direct and practical guidance any court has given on this issue.

Now, what this arc tells us is that the facts drive everything. *Heppner* had every conceivable factor pointing against protection and got the restrictive result. *Gilbarco* had every factor pointing toward protection and said the discovery request was a fishing expedition. *Morgan* is in the middle: work product applies, waiver does not follow automatically from AI use, but the protection has real limits and its enforcement has costs. The practical gap between a corporate defendant with enterprise AI infrastructure and a pro se plaintiff who can only afford consumer tools is now a part of the federal record. It will be cited. And the next court that has to answer the access-to-justice question Magistrate Judge Braswell raised will not be writing on a blank page.

TAKEAWAYS

All right, let's talk about our takeaways. What do we actually do with these three decisions? Here's what I think matters most for litigators and discovery strategists right now.

Think carefully before moving to amend every active protective orders in your cases. I know that sounds counterintuitive, but amending the protective order is a two-way obligation. If you move to add AI restrictions, you are restricting yourself as much as opposing counsel. Before you file that motion, you need to know exactly what AI your client is using, whether those tools can meet that standard, and whether you can actually comply. This is another situation, much like we found with hyperlinked files, where you have to know so much about your client's data before you draft any kind of obligation for them in a protective order or an ESI protocol. *Morgan* gives us useful language from one court, but that language leaves open significant questions: it does not address how data processed through vector embeddings can be technically deleted and verified, it does not address what happens if the AI provider receives a subpoena or government demand implicating data you uploaded, and it does not address the breadth problem I just raised — where drawing the line between AI tools that require contractual protections and the AI features baked into every piece of modern litigation software. Those are open questions you will need to address in your specific protective order language. The language from Magistrate Judge Braswell here should be a guide to you, and it should be something that you continue to refine based on the specifics of your case. Be sure to think them through very carefully. They will have potentially substantial implications.

Think hard about the over-designation problem that Magistrate Judge Braswell explicitly flags. And that means marking more documents confidential than otherwise. She warns the parties directly against over-designating documents as Confidential, and the reason is exactly what you would expect: every document designated Confidential is now a document that the opposing party cannot run through consumer AI. That is leverage. It will be used.It will be used as a battle axe. We are essentially trading one dispute over AI use for another: designation fights. Parties already over-designate constantly. Under a *Morgan*-style order, the incentive to over-designate gets stronger, not weaker, and a pro se or other litigant’s ability to use AI to combat that designation is now gone under *Morgan*.If a document's been designated Confidential Information, they can't input it into the system to ask it how they could refute that confidentiality. Expect that to generate motion practice, and think carefully about your own designation decisions.

Enforcement will be a BIG issue as well. How can you know what a party has done? I'm not sure we've addressed that in the slightest. I want to go back for one second to the motion practice and the confidentiality designations, because we take all the materials that are produced to us and we upload them to platforms. And all that confidential material — whether it's confidential, highly confidential, attorneys’ eyes only, however it's designated — is then subject to that three-step scrutiny that Magistrate Judge Braswell laid out if you have language in your protective order that mirrors what was recorded here in *Morgan*. So be very conscientious about what sort of additional obligations that places on you in those tools, because Relativity has AI review. There are all kinds of tools out there. Relativity was just the first one that came to mind, but there are so many. Almost all the review platforms now have some sort of AI baked into them.

Next takeaway: Understand what platforms your clients are actually using. Not just “are they using AI” — because answer is almost certainly yes. Which platforms, which tier, and what do those accounts’ terms of service actually say about training use, third-party disclosure, and deletion? If your client is using a consumer account — free or low-cost ChatGPT, Claude, or Gemini — with Confidential Information and there is a protective order in place that prohibits that, you may already have a problem. This is a client counseling conversation you need to be having now. And it means that that disconnect that I talk about, where certain people are handling the data and certain people are making the strategic decisions, needs to come together. It's another reason why that disconnect can be a huge problem in litigation if you aren't solving it. The strategic

thought process, the litigators, the people who are drafting the protective orders, need to be talking with the folks who are handling the data. They all need to be one team together, understanding these legal issues and being able to flag things for each other to ensure they're on the same page. It is a critical piece of modern litigation with ESI.

For pro se litigants specifically, or any type of client of modest means, the options under a *Morgan*-type restriction are limited and the costs associated with it are real. Enterprise-tier accounts that might qualify — ChatGPT Team or Enterprise, Claude for Enterprise, Microsoft 365 Copilot with an appropriate Data Processing Addendum — run roughly $25 to $30 per month per user at a minimum. They require organizational billing structures and are designed for institutional procurement. Now, a pro se plaintiff who needs these contractual protections to use AI on confidential materials may simply not be able to get them. The practical reality is that *Morgan*'s protective order may mean a pro se litigant cannot use AI on the most critical materials in their case, while the corporate defendant's counsel has a full enterprise AI stack. That is the gap Magistrate Judge Braswell identified in her footnote, and it does not have a solution yet, but there's an opportunity for the AI companies to step in and provide a cost-effective secure solution for pro se litigants or agencies and lawyers involved in access to justice. Please feel free to share that suggestion with the AI companies. It would be easy for them to provide that documentation that's required under the kind of protective order that Magistrate Judge Braswell has laid out here.

Next, on work product — remember that your prompts are your most protected materials, and protect them accordingly. The prompts you write to an AI system — the directions you give, the questions you ask about the evidence — are your mental impressions articulated in real time. They are almost certainly opinion work product if they are done at the direction of counsel, or you are counsel, or you are pro se acting as counsel. The outputs are probably qualified work product. The tool name is a harder case, and *Morgan* shows you need specific facts and argument to protect it. Do not make conclusory assertions. And if you represent a client who is using AI without your direction, address that now before the materials become an issue.

If you are pushing to discover your opponent's AI use: make sure you have actual evidence of a problem, not speculation, and make sure your request is tailored to the actual concern. If the concern is that Confidential Information was uploaded to an insecure platform, ask for the account documentation and terms — not a broad disclosure of all AI interactions. A fishing expedition for someone's internal thought process is exactly what *Gilbarco* and *Morgan* say courts will not permit.

I'm still not sure what triggered the actual request for the amended language of the protective order here by V2X. There's no discussion of that in the decision, and there's no mention by either party of exactly why that was even required. So again, I would go back to Morgan contesting the actual amendment to the protective order here.

The larger takeaway from all three decisions is the one District Judge Rodriguez raised at the UF eDiscovery Conference: the law is doing its best to apply existing doctrine to a technology it was not designed for, and the courts are being thoughtful about it. But we need guidance from the Rules Committee and ultimately from Congress on standards for AI use in litigation. These three decisions are the beginning of that conversation, not the end of it.

CLOSE

All right, that's our Case of the Week for this week. We'll have a link to the decision in Minerva26, along with the *Heppner* and *Gilbarco* decisions, and other cases cited herein. You can use the Generative AI issue tag to be notified the moment a new decision drops on this issue. If you have questions or want to dig deeper, bring them to the next Meet and Confer session or drop them in the comments.

Please share this episode with your friends and colleagues – staying up on these issues is critical. It's timeliness. It's very, very important and it's hard to do, so we all need help. Subscribe to our blog at Minerva26 to follow all of our Case of the Week updates. If it's about the discovery of ESI, it's in Minerva26, your discovery strategy platform. Thanks for being here.