

# Drawing Conclusions

## Negotiating AI in Your ESI Protocols

A short story by Relativity

Illustrated by designer Kael Rose, working alongside generative AI



Litigation teams aiming to incorporate generative AI in their discovery workflows sometimes face an uphill battle in explaining and defending the technology. What might opposing counsel say? What do cooperation and proportionality look like in practice?


The cast of our story—inspired by real practitioners in the legal community who played out this mock argument at conferences for rooms full of their peers—aims to shine a light on the unknown.

**Let's take a look at how case teams can effectively advocate for the efficiency, accuracy, and defensibility of AI for e-discovery.**





**Janine Delarosa**  
Partner, Acme Plaintiff Firm



**Aurora Bryant**  
Senior Counsel for eLitigation,  
Civil Rights Division



**Simone Davis**  
News Reporter, WXYZ-TV




**Brenda Miller**  
Disgruntled Customer



**Joseph Teegan**  
Discovery Counsel,  
Buzz Nectar Syndicate



**Keith Beckett**  
Partner, Beckett & Fiction LLP




**Cristin Traylor**  
Senior Director, AI Transformation  
& Law Firm Strategy, Relativity




**Mike Markman**  
Cameraman, WXYZ-TV

Meet the  
**C**AST!

**N**ot so long ago, in a place not so far away, there unfolded a matter much like the one you might be working on today...



Buzz Nectar Syndicate (BNS) may be in trouble. After they began sourcing the CBD for their B Serene Honey line from a new manufacturer, distressed customers began reporting adverse side effects—including hallucinations.



Now, a class action suit and related Department of Justice investigation allege BNS knowingly began formulating their products with THC without properly updating labels or warnings.




After brief negotiations, the parties have agreed on almost all of the provisions in the ESI protocol proposed by BNS counsel. But the DOJ had pointed out a detail they've never seen before—and they want to discuss it further.



Counsel for BNS, with help from their technical expert, explain their plan: use AI to accelerate document review, while maintaining transparency through citations, rationales, and statistical validation—just like they would with court-approved technology-assisted review tools.



The DOJ remains skeptical. Their counsel raises questions about the trustworthiness of the tech. Surely it's too new, and too unproven, to be used on a real case?



But the technical expert for BNS has answers. He explains that the AI has been tested thoroughly, and found to deliver even greater accuracy than traditional, manual human review—long considered the gold standard in e-discovery.



In the interest of transparency and validation, the DOJ requests that BNS disclose their validation metrics so they can verify accuracy, too.




BNS, eager to collaborate and get this project across the finish line, accedes to this request. Hesitantly. "So, are we ready to agree on the protocol?"



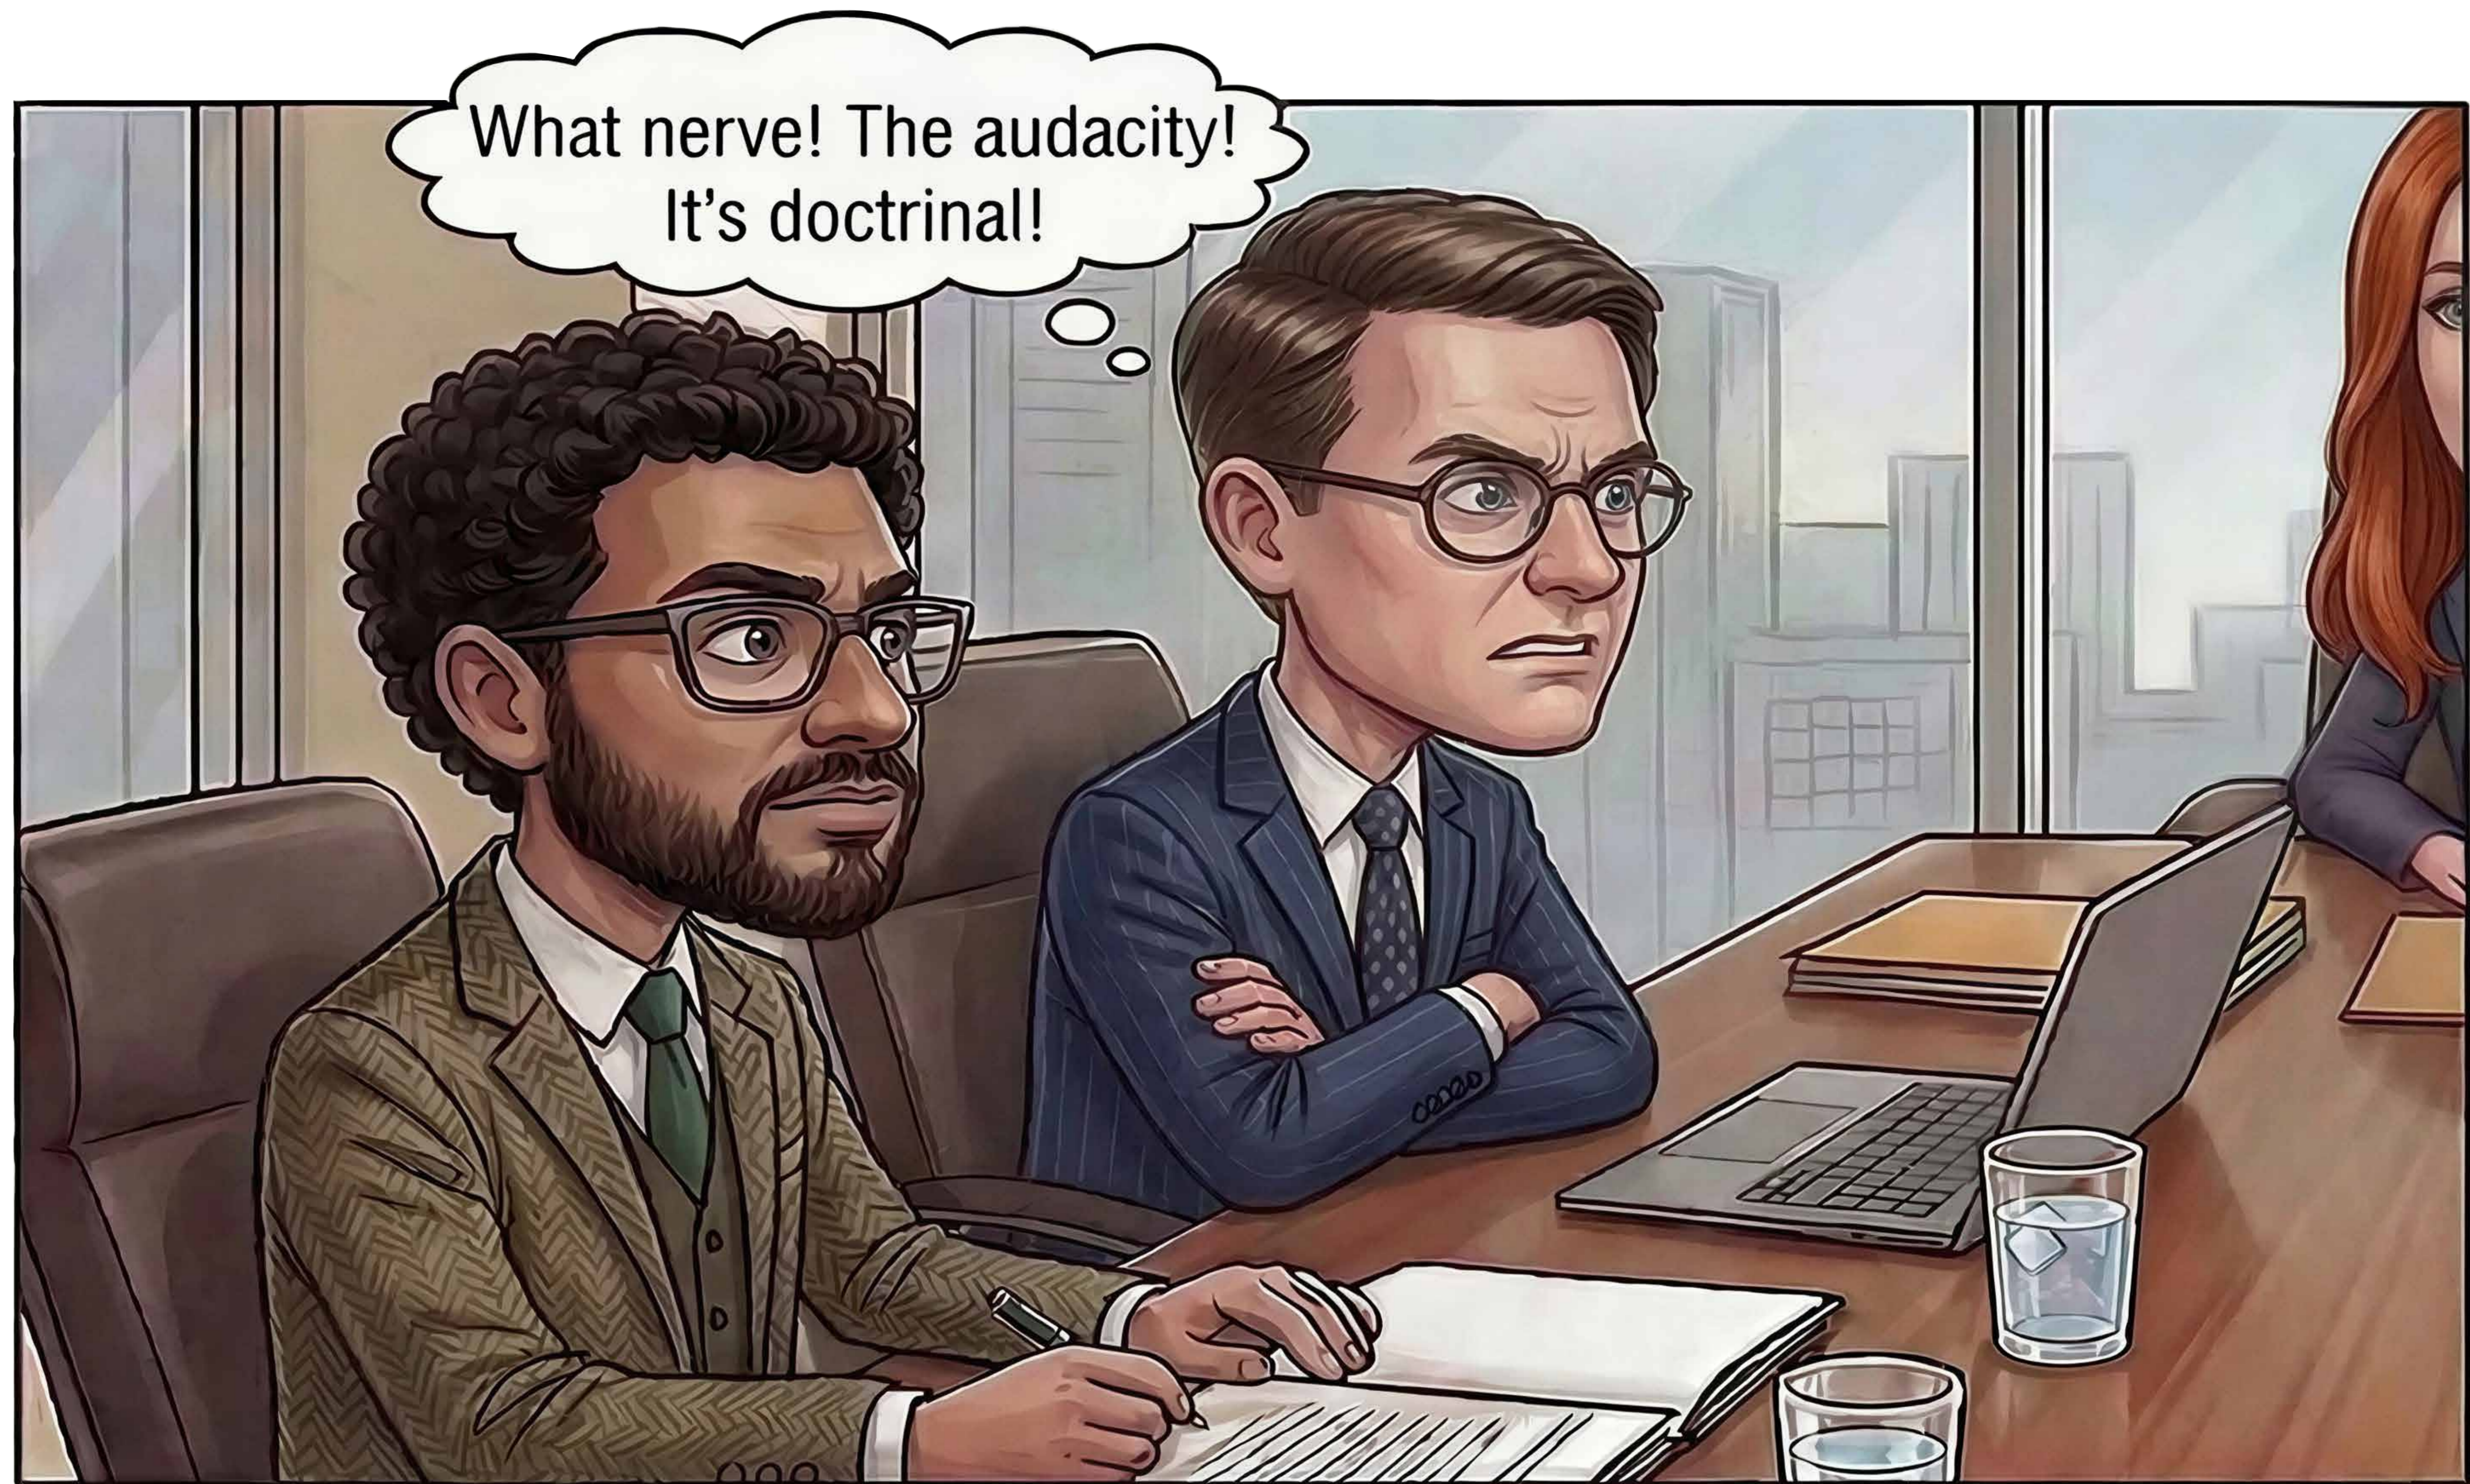
Plaintiffs' counsel interjects, knowing that generative AI requires prompts. She's heard the rule "garbage in, garbage out," and wants reassurance that defense counsel is prompting and training their AI appropriately for the case.



The DOJ is immediately agreeable. They, too, would like to know how BNS plans to train and prompt their gen AI to get accurate results. Both counsel request that BNS disclose their prompt criteria as part of the protocol.



The head of litigation for BNS is quick to correct the course before discussions head off in an extreme direction. "First, you should know the models are pre-trained and we aren't inventing any wheels here. It's fit-for-purpose AI for document review, and our prompt criteria direct it on how to categorize our data set—just like a reviewer protocol directs human reviewers."



He's immediately supported by his co-counsel, Keith, who cites the work product doctrine and declines the request outright. "Our prompt criteria represent my legal interpretation of what this matter is about and how I guide our review. That's protected and I'm not comfortable sharing it with you."



A debate ensues. Is it work product? Or are prompt criteria more like search terms, which are frequently exchanged during e-discovery as a standard workflow to ensure alignment between parties on what defines responsiveness?



Back and forth the parties go. Can plaintiff's counsel suggest prompt criteria? Shouldn't BNS counsels' certification of their production be enough under the ethical and professional rules? Why not use the legalese straight from the RFP as the prompt? Can rationales and citations be disclosed? What about non-responsive document details? Is the AI secure and protective of data privacy?



After much debate, the parties finally reach an agreement: suggested prompt criteria from plaintiffs and the DOJ will be used by defense counsel to guide the creation of their own prompt criteria, and validation metrics will be shared.



Counsel for BNS congratulates and thanks the room. "I'm glad we got this wrapped up. Now, let's focus on what's really important: the issues of the case. And, maybe, finding a nearby happy hour?"

## Epilogue

What should you be prepared to discuss when you're negotiating an ESI protocol and want to include the use of AI?

### Transparency and Trust

- How does the tool guard against hallucinations?
- How can you assess accuracy at scale?
- Is it fit for purpose and securely built and hosted?

### Validation Methodology

- What statistics can you calculate to build confidence in your results?
- How do you plan to provide this information to opposing counsel?

### Workflows and Expertise

- What steps are involved in the workflow?
- How do you keep humans in the loop?
- What model does your platform use?
- How was it trained?
- Is data privacy maintained?

### Cooperation without Waiving Privilege and Other Protections

- How can you reassure opposing counsel of the quality of your prompts?
- How can you provide transparency around validation and QC processes without disclosing your client's private information?
- What culling methods will you use before running matter data through the AI? How will exceptions be addressed?



**Relativity**

[www.relativity.com](http://www.relativity.com)