Minerva26

Carvana Hyperlinked Files: When “Contemporaneous” Production Is Required

Introduction

Welcome back to the Meet and Confer podcast—and to our first Case of the Week segment of 2026. If this is your first time joining us here, this segment does a deep dive into one specific discovery ruling and turns it into strategy that you can use immediately. My name is Kelly Twigger, I’m the Principal at ESI Attorneys and the Founder of Minerva26. Thanks for joining me today.

Today’s case is a perfect example of why staying current matters: courts are not letting parties get away with “the cloud is complicated” when the evidence lives behind hyperlinks in the form of hyperlinked files that may be relevant.

Discovery case law moves faster than most legal teams’ muscle memory. And that knowledge gap is where sanctions happen, budgets blow up, and credibility can get lost. Following discovery decisions isn’t just about “professional development”—it’s how you stay defensible in the constantly changing world of ESI. Because judges are telling us, week after week, what they expect parties to do: what “reasonable” looks like, what “proportional” really means in practice, and what excuses they’re no longer buying. If you’re not tracking those signals, you’re likely negotiating ESI protocols like it’s 2012—or not at all in some cases—litigating discovery on autopilot, and possibly handing the other side leverage that you really didn’t have to give away.

This week’s decision is hot off the press – published yesterday in the matter titled United Ass'n Nat'l Pension Fund v. Carvana Co. As I mentioned, this is a decision from January 12, 2026 from United States Magistrate Judge John Boyle in the District of Arizona.

Now, before I get into the January 12, 2026 order, I want to give you the clean context—because this isn’t a one-off decision, as a lot of them happen in the discovery world.

This decision is built on an ESI order as well as two earlier rulings in the case. First, the August 21, 2025 order, that's the first ruling following the ESI order. That's the order that stated the hyperlinked documents framework that the Court outlined as to how the parties were to proceed providing hyperlinked documents and it’s the order that plaintiffs are trying to enforce here. We covered that August 2025 decision on our eDiscovery Day webinar on collaboration tools, because it’s a perfect example of how Google Workspace and modern collaboration change discovery obligations.

In that August 21, 2025 order, the Court refused to excuse the Defendants from producing contemporaneous Google Drive hyperlinked documents simply because they chose a cloud workflow, and instead the Court ordered a limited test—Plaintiffs could pick two custodians, and Defendants had to use reasonable efforts to produce versions as closely contemporaneous to, but preceding the email communications, as feasible. So basically, whatever technological capabilities they had to produce contemporaneous versions of hyperlinked files for those two custodians.

Now, the discussion at the time was the use of the tool FEC, Forensic Email Collector, which we've covered a number of times on the Case of the Week and I've done a specific podcast on Meet and Confer with Arman Gungor, the CEO of Metaspike and creator of Forensic Email Collector. So we've gone through the technical capabilities of that tool and what they're building it out to be. In this particular case, in August, the Court ordered that plaintiffs could pick two custodians and Defendants had to use reasonable efforts, whether FEC or otherwise, to provide contemporaneous versions.

The second order was on a separate topic, and that expanded the custodians to 25, so instead of 2 we're now at 25 custodians. That matters because the January 2026 order lets plaintiffs select emails from any of the 25 custodians—and that number doesn't come out of nowhere. It's tied to the custodian dispute that the Court resolved in November.

With that context for what we are discussing, now let's talk a little bit about why hyperlinked files matter: hyperlinked files are now evidence, not "nice-to-have in context." There are multiple decisions that we've discussed, more than 40, laying out that hyperlinked files, when relevant, are evidence. If your opponent can't—or won't—produce what the hyperlink pointed to as of the time the email was sent, then your timeline can get fuzzy, your "who knew what when" story can get mushy, and the case can become a credibility fight instead of a facts fight. And in 2026, looking at this decision here from Judge Boyle, the courts are done pretending hyperlinks are somebody else's problem.

What's critical to understand on the issue of hyperlinked files is that the systems that we use to create, store, send and receive ESI are not designed for us to go and get it in a way that works for litigation. That's why we've built collection tools and review platforms to help us collect and review information. But every time we have a new technology it requires us to build new tools. And as that technology keeps changing, the tools have to keep changing. And there has to be ROI for those technology companies to build those tools to allow us to collect that information. It's a constantly moving target.

Here, the use of hyperlinks to share documents has created a new technological challenge that is not yet solved – and that is the ability to collect the copy of a file from a hyperlink that accompanies the so-called "parent" communication, or the originating communication. These days, I can share a link to a document in a text message, instant message, email, any way you can communicate electronically. And there are two issues with getting that file – getting access to it so you can collect it at all, AND whether you can get the version that existed at the time the communication was sent – the contemporaneous version. Systems that do let you collect will

give you the most CURRENT version, generally. But what existed at the time the communication was made may be key for the case. And that's what we are dealing with here – when a party is required to go and get the contemporaneous version of a hyperlinked file for discovery.

Let's dive in.

Facts

The underlying litigation here is a securities fraud class action in which the Plaintiffs allege that Carvana and certain executives made materially false and misleading statements that inflated Carvana's stock price, in violation of the Securities Exchange Act.

As I mentioned, this one is before Magistrate Judge John Z. Boyle in the U.S. District Court for the District of Arizona.

Now, this January 12th decision ruling from yesterday that we're talking about, this is a Motion to Compel Compliance with the Court's earlier order in August—meaning that we're not debating discovery theory anymore. We're in the "do what the Court already told you to do" phase. And the order that we're focused on is from yesterday, that's enforcing the August 2025 framework.

When you are talking about hyperlinked files, the platform you are using really determines the technological capabilities with what you can do. Carvana uses Google Workspace and Google Vault, the archive storage space for Google Workspace. In Google Workspace, people don't attach files the way that they used to. They drop a link to a Drive document: a spreadsheet, a deck, a doc, etc. The email contains the link, but the content lives somewhere else. And here's the problem: the Drive document can change. It can be edited. It can be overwritten. It can be shared, re-shared, renamed, moved, deleted, restored—welcome to collaboration culture.All of those things can happen in the time between when the link was sent in the originating communication you're interested in and when you're trying to collect the information. And so the question is how do you get back to that point in time that we are used to having with physical attachments?

So the Plaintiffs want the version of the linked file that reflects what the recipient would have seen or known at the time the email was sent—what the Court describes here as the contemporaneous or point-in-time version. Defendants, on the other hand, are arguing feasibility, burden, and the limitations of getting accurate "as of" versions inside Google's ecosystem—especially at scale, across custodians, over time, with version histories and access rights and the rest of the chaos that exists in the cloud environment.All of those things could likely be relevant, but it's the feasibility of getting them technologically balanced with proportionality considerations that are always at issue when we're talking about hyperlinked files and the information that surrounds them.

A little bit of background here from the August 21, 2025 order. That one the Court looked at the original ESI framework in which the parties had agreed to provide hyperlinked documents. And that order also required the parties to use reasonable best efforts to collect and produce

hyperlinked documents where feasible, and to work collaboratively on things like preserving the hyperlink relationship and producing point-in-time versions.

The parties then collided over how to do it. Plaintiffs wanted contemporaneous versions of hyperlinked Google Drive documents tied to emails—because they said they needed it to establish “who, what, where and when.” The Defendants said that wasn’t reliably doable, and they attacked plaintiffs’ proposed tool—Forensic Email Collector, or FEC—as inaccurate, burdensome, and producing documents that couldn’t be authenticated.

The Court didn’t really let either side win the argument in the abstract. Instead, Magistrate Judge Boyle ordered a limited test: plaintiffs could select up to two custodians, and the Defendants had to produce responsive hyperlinked documents by December 1st. Those had to be as close to the contemporaneous version as possible. The Court assumed that FEC would be useful, but left room for other agreed methods. And the Court basically said: if the test works without undue burden, the parties can come back and talk about expanding.

Now the November 19, 2025 order is a completely separate issue altogether, but what’s important is that the Court expanded the custodians from two to 25. We mentioned that a little bit earlier, but that broadens the landscape for the hyperlinked files issue here.

So the current landscape that we’re in in January 2026 is this agreement to provide hyperlinked files, the testing process the parties are supposed to go through, and now we’re talking about 25 custodians.

Plaintiffs filed a new Motion to Compel Compliance saying, in substance, that Defendants weren’t complying with the August 2025 order because they were severing hyperlinked documents from parent emails, and resisting the production of meaningful point-in-time versions.The Defendant's arguments are pretty much what we've discussed, which is that it's too burdensome, technologically, we can't do it at scale.

In the Court’s discussion, Judge Boyle boils down the fight (no pun intended there) to two questions: how far Carvana has to go using Forensic Email Collector—FEC—to retrieve the most contemporaneous versions of hyperlinked files across all 25 custodians, and whether Carvana has to produce responsive, non-privileged documents sitting in Google Vault. And the Court takes a moment to level-set on the technology: hyperlinks aren’t “attachments” in the old-school sense because they’re not fixed in time, but that doesn’t make them an option. The Court recognized that Google Drive and Vault create real production challenges, but it’s not going to let novelty become a loophole that blocks discovery—especially where the point of the request is timeline evidence: who knew what, when.

This is consistent with other rulings that we've seen in the past in the In re Uber Techs. case, in In re Meta cases. So this is consistent with what we're seeing from judges. They're no longer looking at Google Vault and saying, if you chose that as your storage capability, then you're stuck with the burdens that that imposes upon you to provide hyperlinked files in discovery.

On the FEC question, whether or not that tool should be used, the Court anchored itself in Rule 34's "reasonably usable form" requirement and the idea that ESI shouldn't be produced in a way that makes it unreasonably difficult for the requesting party to use. Then it struck a benefit/burden balance by building a bounded, scalable mechanism: the Plaintiffs got to choose 250 responsive emails from any of the 25 custodians, and the Defendants must run FEC to retrieve the most contemporaneous versions of any non-privileged hyperlinked documents tied to those emails—without limiting this to the two custodians from the earlier test. And importantly, the Court made clear that the difficulty of point-in-time retrieval didn't justify withholding linked documents altogether; if the "most contemporaneous" version is hard to pin down, that's a reason for structure and limits—not a reason to say "no."

On Google Vault, the Court separated two ideas. It found that producing responsive, non-privileged Vault documents was proportional and consistent with the ESI order, and it expected those documents to be produced. But it declined—at least for now—to require a full-blown "reverse search" that the Plaintiffs wanted from Vault going back into custodians' mailboxes to locate parent emails, because the Plaintiffs didn't show how that would work or why it would be proportional and not unduly burdensome. Instead, the Court went with a capped, targeted approach: once Vault production is made, Plaintiffs can select 200 documents and Defendants must use traditional ESI techniques to try to locate a parent email if one exists, and produce it if it's not privileged.

The Court's solution here is very practical. It's not "produce everything forever in every version." And it's not "if it's too hard, don't worry about it." Instead, it's essentially a sampling approach that we've seen done in a number of contexts over the years with different electronic evidence sources. And that is that it has bounds around it, it's a phased approach, and it's testable. And the Court started with that in August and is following up on that in January because the Court is trying to get to the truth of the facts in the case without letting process become the case. And that's where we break down often in discovery with ESI is that we're letting the process, the compliance with the process, drive what we're doing instead of backing up and thinking strategically about where we are in the case and what the facts are that we want to gain.

Here's what Judge Boyle ordered:

An email-driven, capped selection in terms of the emails that were produced: Plaintiffs may select 250 responsive emails from any 25 ESI custodians. Defendants have to then use FEC to produce the most contemporaneous version of those within 10 days. So, an actual process defined.

For Vault collections, Defendants must provide all responsive, non-privileged documents from custodians' Google Vault collections. Plaintiffs can select 200 documents from that set, and Defendants must search for a potential parent email for those documents—and if they find it, then they have to produce it. So emails and documents separate because in Google they're stored separately. You have Google Mail and you have Google Drive, which stores documents. So two different places that you have to provide for.

In essence, the Court took the August 2025 pilot concept, layered in the expanded custodian universe from November 2025, and built a structured production pipeline—because “cloud is hard” doesn’t excuse missing evidence when the links are actually the evidence.

Takeaways

First: Write hyperlink rules into your ESI protocol upfront. Define what counts as a hyperlinked “attachment,” what must be collected, and whether production must be point-in-time/most contemporaneous version versus “current version,” including required metadata (file path/ID, timestamps, version indicators, sharing context where relevant).

Second: Run a defensible pilot and document it like you mean it. Pick a capped sample (like the Court did), define success criteria (accuracy, completeness, turnaround time), record the steps/tools used, and memorialize the results—because “we tested it and here’s what happened” beats “we don’t think it’s feasible” every day of the week. A court is always going to look at exactly what you've tested and what you know your capabilities are in ruling on these potential technological challenges.

Next: Argue proportionality with evidence, not with vibes. We've talked about this. And so I was trying to think about over the holiday break what's the best way to get this across. But we see so many times counsel just saying: “Well, we feel like this is the best way to do it, judge. This is gonna be very difficult. This is gonna be hard. You've got to have evidence. You can't just state how it is that you think things are gonna go down. Courts are requiring concrete facts about what you're seeing technologically in order to make these arguments. If you’re resisting, quantify the burden in terms of time, cost, error rates, technical constraints and propose some alternatives (sampling, a phased approach, narrowing custodians, prioritizing email sets). If you’re requesting, tie the ask to the case story: timeline, knowledge, reliance, intent—why this evidence matters.

Next: Treat parent-email linkage as a core relationship—not a bonus feature. If documents are coming from Vault/collections without their transmitting email, build a process to locate and produce the parent email (or explain, with proof, why it can’t be found). Don’t let your collection method destroy context and then call it “reasonable.”

Next: Don’t walk into a meet-and-confer with “cloud is hard.” Walk in with a plan: tool options, workflow proposals, limits you can live with, and a timeline. Judges want to see collaboration and concrete proposals—especially when the Court has already issued an order and is now managing execution of that order, as we see here.

Manage privilege deliberately for hyperlinked files. Decide how you’ll handle privileged content inside linked documents (comments, tracked changes, embedded materials), what gets withheld versus redacted, and how you’ll log it. Hyperlinks multiply the privilege risk because they multiply versions and collaborators.

Finally: Remember the credibility tax. Sloppy or evasive hyperlink production invites motions, court oversight, and the kind of skepticism that bleeds into merits. If the “as-of” version matters to “who knew what and when,” treat it like real evidence—because the Court will.

That’s our Case of the Week segment for today—and if you take nothing else with you, take this: hyperlinks are evidence, and courts expect you to treat them like evidence. If your protocol doesn’t address point-in-time versions and linkage, you’re not “being modern”—you’re building risk into your case.

This is also why we built Minerva26—the discovery strategy platform that connects case law, rules, and real-world workflows—so teams can turn decisions like this into a defensible plan instead of reinventing the wheel in every meet-and-confer.

If this episode was helpful, please do me a favor: share it with a colleague who’s wrestling with Google Workspace, Drive links, collaboration tools, or other discovery issues, and post it on LinkedIn with your biggest takeaway—because the fastest way to level up discovery strategy is to normalize these conversations among counsel. And if you haven’t already, please subscribe to the Meet and Confer podcast so you don’t miss the next ruling that changes what “reasonable” looks like.

Thanks for listening—and we’ll be back next time with another episode of Case of the Week.