Mathieu (00:27.05)
Well, I dressed up for you today, Eric. So I'm looking forward to this conversation. I've had the opportunity to hear you present a couple of times about your, I guess, different initiatives. It might be worth it just at the start to talk a little bit about the, I know there's different working groups, there's a standards group, there's another group. So maybe just to differentiate these things for the listeners. before we get into the conversation. And maybe as you're doing that, talk a little bit about the problems, the statement that led to these things being formed.

Eric Scouten (01:05.05)
Certainly, of course. So thank you for having me. Yeah, I think it's a good question of what are the different groups that are involved here? And I can talk as a member of and to some extent representative of a few different organizations. The biggest one is the Content Authenticity Initiative itself, which was formed as a business partnership with Adobe and a few other companies at the end of 2019 with the shared goal of addressing the concerns about misinformation, disinformation, misattribution of information that was designed to mislead and misrepresent content on the internet. And that led to, you know, a pretty wide membership of organizations that are sharing that purpose of understanding how misinformation exists and what we can do to differentiate legitimate information from that which is designed to mislead.

Eric Scouten (02:10.00)
Separate from that, there's a group called the Coalition for Content Provenance and Authenticity, often abbreviated as C2PA, which is a technical standards organization formed within the Linux Foundation as a joint development foundation. And that publishes the definition of how you assemble these pieces of information, how you represent them in binary format, and how you embed them or host them outside of an asset if you want to put them on on some kind of cloud storage or something like that instead.

Eric Scouten (02:40.14)
And I want to circle back to Content Authenticity Initiative. That is also the name of the business department within Adobe, of which I'm a part, which manages our interaction with these two organizations, but also manages integrations of this work that we do at open source and open standards into Adobe products as well.

Eric Scouten (03:08.17)
And then finally, there's a new group that has just started up recently called the Creator Assertions Working Group. And that is a separate standards organization from the C2PA, but is designed to augment the C2PA standards with additional data types that are particularly relevant to human and or individual and organizational identity as it pertains to the C2PA ecosystem.

Mathieu (03:32.10)
It seems like describing, at least from, not that I struggle with it, but I find it sometimes difficult to talk to people. You can have family members that don't really understand what I do for a living, what types of things I'm building in the digital trust space. But it seems like explaining content

authenticity to people is much, much easier, I guess, just because people understand this problem or face this problem in their everyday lives. And I'm seeing more and more utility for this type of thing. Like, I guess all of this work was started before the acceleration of generative AI that happened just maybe over a year ago, if I just use the start of 2023 when chatGPT really got more popular and these, these models became more popular, but it's becoming easier and easier to just spoof any type of media. So has that accelerated the interest in this and has that changed kind of the vision for any of this?

Eric Scouten (04:36.02)
Um, I would say that it's definitely accelerated the interest and adoption of what we do. Um, that said, uh, you know, the, the pitch that was made to me to join this project at the end of 2019 included a prediction that generative AI would be disruptive to the marketplace of ideas. And, you know, I think the only thing that we've been surprised by is how quickly that's happened.

Um, but we kind of knew that it was coming. And of course we have our own generative AI tools as well at Adobe. So we understood this to be part of the problem space. And the thing that we have really concentrated on and the thing that resonates to me when I talk with friends and family and people that are outside of this workspace, if you will, is that people understand that they're having a harder and harder time discriminating between legitimately constructed content and that which is designed to amuse or mislead or what have you. And so the analogy that I often make is the transition that we made some years ago between unsecured internet traffic and HTTPS. You know, I happened to be looking back at some content I wrote a decade or so ago. And I was stunned by the number of plain HTTP links that I referenced at that time.

Eric Scouten (06:03.07)
And today you would look kind of askance at that and sort of ask questions about, well, who's potentially changing the content on the way from the intended server to my browser? And our hope is that we establish a similar level of sort of introspection and question asking as people view images, media, audio recordings, so forth and ask that same kind of question, well, do I know how it got to me? **And we're providing tools that help people answer those questions**. And so there's sort of a virtuous cycle of people being more and more interested in understanding the provenance of their content and news media and other organizations and other sectors being more and more interested in providing that, that authenticity, that provenance.

Mathieu (06:59.13)
Might be a dumb question, but one of them, I was actually surprised recently that one of the financial institutions that I deal with was trying to push me towards using a voice authentication within their telephone system so that when I call, I could actually get authenticated quicker. Just knowing how easy it is to replicate these things today, I found it quite odd that in 2024, these types of authentication solutions are being pushed forward.

All this content authenticity initiatives are, and again, it's maybe a dumb question, but are you looking to solve problems of existing content or is there also a scope that encompasses like live

content? And like we've seen, we see examples of people getting defrauded on the telephone and stuff like that. Is that in scope or is that completely out of scope?

Eric Scouten (07:53.03)
It's a question that's come up several times recently, and I've spoken to people who work actually specifically in banking in various countries. And there's definitely a concern about, how does the bank differentiate themselves from those who would like to steal your banking credentials? And this is a variation on that, how does the bank understand who you are? That said, the specific case of voice communications over the internet is not something that we have treated as particularly in scope. Not so much because we think it's not a problem worth solving, but because the people who have stepped forward and been interested in this ecosystem so far really have more expertise in broadcast use cases. So as a BBC or a New York Times or other major organization such as that with a large audience wants to make sure that their content is differentiated from those who would appropriate their identity. That's something we're very much prepared to address and it's in scope for us. The question of communication, I think there are adjacent technologies that are aiming to discuss that, but I don't know that I've seen a direct case of that overlaps with the approaches that we're taking in C2PA.

Eric Scouten (09:19.16)
One thing I will say is that in the C2PA space, there is quite a bit of attention being paid to streaming media formats. So you could anticipate that a live player for a broadcast audience could annotate with notations about how the stream is being constructed on the fly. And there are some prototypes that I have seen of that are pretty promising.

Mathieu (09:43.04)
Who needs to be involved within that, if I call it supply chain for that provenance chain of trust, however you want to call it, to be captured? I know there's hardware companies or just device companies that are part of the initiative, but is looking at this like just a traditional supply chain provenance problem to be solved the right way of looking at it?

Eric Scouten (10:28.23)
I think that's a pretty close analogy. And we think about when we present on this, we think about that supply chain of capture device, of editing workflows, of publication workflows, of perhaps re-rendering workflows, such as a CDN or similar might do. And then the viewer being able to introspect, hopefully all stages of that. But that really requires all of those stages along the way to understand the standard and to be willing to do that instrumentation. And there are growing adoptions in all of those sectors, but each of those has its own sort of piece to unlock.

Eric Scouten (10:49.22)
You mentioned cameras. There are, Leica has already come out with a camera that's in the market. And I believe two other camera manufacturers have announced plans to be in the market during this year with devices that capture directly on camera. And then there are smartphone applications that can do that outside of the standard camera applications. And then

I think with the recent announcement that Google has joined the C2PA, you might expect that they would put that into some of the devices that they manufacture as well.

Mathieu (11:23.11)
Does most content that gets created and then distributed on the internet come from mobile devices? Are there a small amount of types of devices or devices where you would think it would be easy enough to at least have... You need a starting point in that supply chain. Capturing provenance within a supply chain can get very complicated very quickly depending on the type of use case or the amount of participants involved. So is it easy enough to pinpoint like, here's a large percentage of where the supply chain starts and then maybe pinpoint where there may be some low hanging fruits or easier supply chains to capture with this?

Eric Scouten (12:08.17)
One of the challenges I think that we face in answering that question is that we don't necessarily have insight into the universe of content that is published. Right. So, you know, the place where I think I would start to answer that is if you're looking at a news media organization who is motivated to be part of that supply chain, what do they ask the people that are providing content to them to provide?

Eric Scouten (12:32.11)
So you know, you could imagine an AP or a BBC or something like that saying to their reporters, you need to have devices that produce this kind of metadata in order for us to then publish it through that workflow. That said, the standard was written with the idea that not everything starts on day one. And so, you know, there are provisions for documenting what content looks like when you brought it in if it doesn't have C2P metadata on it. Um, and you know, would you love to have a complete supply chain? Yes. But you can also document the fact that you don't and take responsibility for, I brought this in from this source that is not verified in the same way. From here forward. Here's what I'm attesting about that, that piece of content.

Eric Scouten (13:25.12)
And you can imagine, for instance, news organization, if they're still publishing a photo taken in 1969 or the moon landing or whatever, obviously there's no C2PA in those days. There's not even digital photography for the most part. So they might still be able to attest, I have this and I know from some other source, some other mechanism that it was taken on this date in this place and I'm attesting to that now, but I can't retroactively attest to that with an original supply chain, if for instance, it was captured on film or older digital cameras.

Mathieu (14:02.00)
Yeah. And I guess like striving for protection is maybe the enemy of having something good or something better than we have today. I think just starting to see signatures on pieces of content kind of elevates the game from where it's at today. So, um, and you would think too that over time, like if certain content supply chains have fuller provenance therefore leading to maybe higher integrity. These pieces of content may be more valuable than others and it could create some kind of headwind for others to want to produce content that, you know, if I'm looking at this

one and they're showing me the provenance and this one's not, then the one that isn't maybe will want to do so and adopt this to at least showcase the integrity compared to the other one.

Eric Scouten (14:51.06)
Yeah, I like the concept of headwind. I think that fits very well with what we are thinking about.

Mathieu (14:57.07)
What are the different types of content that fit into here? We've talked about like images or videos or streaming, like would news articles fall in here? like is there a full list of content types and then does the way we look at metadata across these pieces of content to differ?

Eric Scouten (15:15.15)
So the ones that I know of that are formally specified are still images, videos, audio recordings, PDF and similar documents and font files. And there's a provision that for something where it isn't specified, you can do a sidecar that says, I, you know, attest to this history of this, this set of binary data that is encompassed in this other file adjacent to it. So for instance, plain text files, since you can't really embed anything in that, you could do a sidecar for a plain text file that is a C2BA manifest. But, and I think most of the major publication formats for each of those classes of media that I've talked about are already well described.

Eric Scouten (16:05.02)
It is a question that we're often asked about news articles in particular. I don't think we've found a way to describe that specifically. And part of the challenge there is that the rendering of an article into HTML, for example, could change from day to day as the news organization updates their content management system and chooses new styling or, you know responds to new standards in the HTML world, and that would invalidate the signatures from prior versions. So I don't know, I don't know that that's understood in a way that we're prepared to standardize, but I know that there's definitely interest in that.

Mathieu (16:48.09)
So if a New York Times or a writers or something like that wants to use this, it would mostly be in the short term for like photo or video assets, you would think.

Eric Scouten (16:58.16)
That's the currently supported use case, yeah. And part of the situation is, content that they publish on their sites often gets taken and either miscontextualized or reformed into something else that might still have a New York Times logo on it, but now be used in a way that misrepresents what the New York Times originally said about that surrounding context.

Eric Scouten (17:28.21)
And so what we enable the New York Times or any news media organization to do in that case is to make the statement about, well, this happened at this location at this time and here's what we reported about that situation at the time, and any effort to say anything else not signed by us, probably not realistic.

Mathieu (17:49.06)
Do you see this being used for user generated content as well? And how does that differ then? Cause if a New York Times is using this, there would be a root of trust, I guess, within the organization, there needs to be some, some list or some governance around, Hey, this is the organization, These are the signing keys type of thing. And then you could even extend, extend that to different journalists or content producers within that media outlet type of thing, but if you're just starting with like end user generated content, would it be as easy for use from just like you and I just creating content ourselves?

Eric Scouten (18:29.02)
I think there's an interesting challenge, and this is, you know, reflecting back to the identity community a bit, is that it's almost there. You know, there are varying forms of digital identity, and depending on what state or country you live in, you may have access to more than, you know, people in other places. And can we take those kinds of either government-backed identity or identity that's backed by, know your customer vendors who look at your driver's license and say, you know, here's the name that I found there, and tie that into this. This is definitely an area that we're very interested in moving things forward so that individuals who aren't affiliated with some larger organization with that infrastructure can make those statements.

Eric Scouten (19:15.21)
That said, we're also very mindful of some hazards associated with that. So you could imagine somebody who is, you know, filming, you know, some kind of demonstration in a country where the government doesn't have due process rights might not want to be identified because you know their presence or their recording of interactions that might have happened there could actually put them at harm, could put them at risk. So everything that we do we do with the idea that it is opt-in, it is redactable at a later point. So if a person or a publisher, a subsequent publisher, deems it a risk to reveal the identity of the original creator, that can be hidden. And we do work with organizations that are representative of citizen journalists in various countries where threats and harms to their safety are a constant issue.

Mathieu (20:16.21)
That's super interesting. Thinking about basically like identifiers that want to be public versus ones that just don't want to be for various reasons, but still, still knowing that whatever authenticity means in that standpoint, but maybe just knowing something is not AI generated and is human generated is.

Eric Scouten (20:39.20)
There's that. Yes, and I think there will be an interesting market going forward in content that is provably not generative AI, in part because generative AI has a problem of starting to get confused when it feeds on its own output. So I think there will be an interesting marketplace for provably human-generated content.

Mathieu (21:01.02)
Why the decision to focus on providence versus detection after the fact?

Eric Scouten (21:06.23)
We looked at the rapid evolution of generative technologies and felt like that was an arms race that we just couldn't win. You figure out what the generative tools are doing this month and two months later they've come up with something that is substantially smarter and outsmarts your detection. So the odds that you could stay ahead sustainably of those improvements in that kind of technology, I think are essentially unwinnable. So we've very explicitly chosen not to be in the market of detection from day one.

Mathieu (21:44.23)
Yeah, so then the first adopters that use this may gain a competitive advantage, which like I was mentioning earlier, may trigger others to want to use this if they want to differentiate themselves from other content that's on the internet.

Eric Scouten (21:58.00)
That's certainly our hope, yes.

Mathieu (22:00.22)
Yeah, it may be worth jumping a little bit now into kind of just more from a technical standpoint, what this looks like, what the C2PA data model looks like. So we talked a bit about different type of metadata, but what does this data model look like? What are the different sections that are in there? And I'm sure I'll have some thoughts or follow up questions from some of these different categories.

Eric Scouten (22:26.00)
Certainly, so I'll start by describing sort of the fundamental building block of history is what we call a C2PA manifest. And that's a binary data structure. We use the jump format, which is basically a way of recursively defining boxes of arbitrary content. And I'll go into what a few of those are as we kind of dive in.

Eric Scouten (22:47.20)
But a C2PA manifest represents one sort of signed action on behalf of one sort of set of edits along the way or acts of creation. And within a C2PA manifest are a handful of things that we call assertions. And those are, many assertion types are defined within the C2PA standard, but others can be defined outside of the standard as well.

And what you do is you combine the set of assertions together into what's called a claim. And that claim is signed by the hardware or software that is producing this overall data structure. And that is the representation of what has happened. Assertions can be any number of different things. Some examples of them are, you know, camera captured location or timestamps of where the original capture occurred.

Eric Scouten (23:42.19)
In the case of Photoshop, we instrumented with actions that were taken during it at your option. So, you know, if you want to say I did, you know, these particular kinds of things, I brought in these other ingredients, things like that, those can be documented. There's some work being done now to document human and individual identity and, or, and organizational identity into what we call an identity assertion that's outside of the C2PA, but in that framework of assertions that can be added, thumbnails, things like that. There are many, many more kinds of assertions that are available. But those are examples of statements you can make about the document or about the asset. And those statements then are gathered together in this claim, which is itself signed. And by doing that, then you provide tamper evidence.

Eric Scouten (24:38.10)
So if somebody subsequently tries to represent something different about that, the signatures fall apart and you no longer have valid C2PA metadata. You can't prevent somebody from changing the bits of a file that's just the world we live in. But what you can do is provide evidence that something like that has happened.

Eric Scouten (24:58.17)
Building a little farther from that, one of the things that is novel about the C2PA approach is that if you have a chain of custody, that entire chain is represented in what we call a C2PA manifest store. So for those who are developers, you might think of Git commit history that builds up over time as new authors and new tools and new work is done on an overall project. And similarly, C2P manifests can be chained to form a arbitrarily long description of the history of the content that you can dig back through and see in various visualization tools.

Mathieu (25:40.04)
In the digital trust space we often talk about trying to provide inputs to people so that they could facilitate them making trust decisions. So the way I look at this is just, it's another digital trust input and people could make their decisions based on that. But does this badge or what's being shown to users on the end experience, does it tell them that something is verified or not or legit or not, or it's really just, Here's kind of a history and up to you to make your mind up based on your context.

Eric Scouten (26:16.20)
Yeah, so we're really careful to say that we're not in the fact checking business or trying to help people say what is true or what is not. It is, as you suggested in the latter part of your question, it's can you verify that this comes from a source that you trust and was it altered after that source produced that media? But the decision whether to follow news media A versus news media B is not something that we want to try to advise you about? And have they reported something accurately? Again, something not that we want to try to advise you about, but more, do you have a history of understanding and believing the content that comes from this organization? And can you continue that relationship by viewing this content? Is that helpful to you?

Eric Scouten (27:06.02)
But you think about the Snopes and the PolitiFacts and so forth of the world that are out there trying to judge the actual truth of a particular story or image or whatnot that's on the internet. We are not in that business at all.

Mathieu (27:20.04)
And for the publisher of this image, they all use the same standard, but is there like a set of tools that they would use that when they're publishing certain media, you almost need to like export your photo through this thing and like, how does it come together? How does it get rendered? How does the website show that? What's the vision and the architecture around that?

Eric Scouten (27:43.06)
Certainly, certainly. So the team at Adobe produces open source code that implements the standard and that's available. I can send you a link for that afterwards that you can put in the show notes. And so, a producer could do their own implementation or they could use the tools that we've provided. We base most of our work in the Rust programming language but there are bridges to an ever increasing set of other languages that can bridge to that Rust code that's the core implementation.

And so, among the tools that we present is a command line tool that you can use. So assuming that you had the signing certificates and so forth in place, you could then use that C2PA tool to add the metadata that's required for a C2PA statement. So that describes the publication side, or you could use a camera or a piece of software that already has that implementation built into it. So, you know, it depends on what tools and what, what tools you already use and how closely are to already implementing the standard.

Eric Scouten (28:54.00)
On the consumption side, you know, the fallback is that we have a site called Verify where you can drag any piece of media that's supported by C2PA, and you can visualize that history. And we have quite extensive set of tools that you can use within that site to drill down into the history of a piece of content, which could be arbitrarily deep. We've tested with histories that have hundreds or I think even thousands of manifests in them. And then that same open source library that I talked about could be used by anybody else to build their own implementation of that and report on whatever aspects of the metadata that they are interested in.

Mathieu (29:40.03)
are the assertions that are made. So the way I'm… if I could just repeat how you described the architecture is that I have an asset, let's say it's an image. And then for my image, I will have a manifest. And then within my manifest, I may have a set of assertions. And if there's a thousand edits that have been done, maybe there's a thousand assertions. And you can view it almost as like a tree that I start with my manifest and then I could have one assertion. I could have another one under it, but maybe there's a couple that split off or come back.

Eric Scouten (30:14.13)
Yeah, so an assertion is a description of some piece of metadata about that act of creation. And it's the edits that we were talking about right now go into a single assertion that's called an actions assertion. And so that actions assertion might have a list of a thousand steps that you took from point A to point B in the production of that image.

Um, unless you were in a multi-author situation where you wanted to document the participation of different people and different parts of that editing workflow, you typically would wind up with one actions assertion, but it could be a fairly large one if, if the history is long. When you start to talk about bringing in other pieces of content, um, to the extent that those pieces of content have their own C2PA metadata, A compliant tool is supposed to document those as ingredients. So the manifest that came with that original content is carried forward into the new one and referenced as a, I placed this image here. Here's its original manifest. So you can see how that tree of different content was produced.

Mathieu (31:30.21)
That's helpful to understand how they would be categorized as action assertions. And of course you could have, if you're blending different assets together, that's helpful to understand as well. And in terms of the signatures then, so there, there could be a mix of human and machine assertions and signatures.

Eric Scouten (31:56.13)
Yeah, so in C2PA 2.0, we made the decision to focus more on, in the core technical specification, we made the decision to focus more on what is machine attestable. So you'll notice actually some parts of the specification were deleted going from 1.4 to 2.0 to honor that distinction. So in 2.0, the outermost signature that is over this claim data structure, that wraps all of the assertions is really intended to be issued by the hardware or software that is implementing the C2PA spec. That was less clear in the 1x versions of the specification. And so that leaves open the question that you've raised about how do individuals and organizations that want to identify themselves with their content, how do they do that?

Eric Scouten (32:50.21)
And that is the work of the separate group that I've talked about, the Creator Assertions Working Group, is really aimed at surveying the landscape of how identity is expressed digitally and how you can bind that to the content. So the expected output of that group will be what we call an identity assertion, which is a separate signature issued by that person described or organization described by a digital credential that also signs over a list of assertions that individual or organization is claiming authorship for. And so you wind up with sort of this matrix of assertions signed by the tool, assertions signed by this credential holder, potentially assertion signed, different sets of assertions signed by multiple credential holders that documents each person's role or organization's role in that context.

Mathieu (33:49.09)
Who manages the public key infrastructure for the hardware or software signers? I would imagine that there would need to be some attribution done somewhere.

Eric Scouten (34:00.20)
Yes, so the C2PA is in the process of spinning up a trust list for implementations of the C2PA standard. So you can expect that, you know, at some point, I think later this year, but I'm not positive of that timeline that, you know, known compliant implementations of the C2PA standard will be Will be able to trace up through CAs that are known to the C2PA and identify as roots of trust for the overall signature. You could also do self-signed signatures that results in a different sort of report from the C2PA specification of a well-formed manifest versus a C2PA valid manifest. That's basically does it anchor up to a known root of trust or not

Eric Scouten (34:51.17)
Separately, and this is a question that the Creator Assertions Group will be taking up, is what would be a similar set of suitable trust routes for humans and organizations to trace through? Because I think that's a very relevant question. I think we're still looking for some answers on that front.

Mathieu (35:11.21)
And obviously, I could talk about that with you all day on the trust registry side of things, but I guess it kind of makes sense now. It's looking at the signature suite and how it works for signing the claims. And I kind of, I think the one that was chosen is very light and well suited for hardware or software versus if you're going to have individuals signing after making assertions that obviously could be just be using different crypto altogether.

Eric Scouten (35:44.02)
Yes, yes, and I think likely will be. There are definitely some organizations that have interest in and familiarity with the X.509 based identities that we basically currently use in the C2PA Core spec. But I also expect that people will want to bring other forms of identity to that space. And in the Creator Assertions Working Group, we're talking through two forms of identity for sure, and possibly a couple others, have been raised as alternatives that we may want to consider as well. If you live in a state that has a mobile driver's license, could you use that, for instance?

Mathieu (36:24.15)
Yeah, it's just where is the root of trust? So if you are talking about individuals and you have a mobile driver's license, that obviously could be a great root of trust if you wanna have to link identity, personal identities with signatures type of thing. And then obviously it seems like there's the whole organizational identity as well where an organization could delegate identifiers or signature to their creators or journalists or stuff like that. So that's maybe more of a federated model.

Mathieu (36:59.10)
Yeah, the whole just any individual creator is going to be an interesting model because then again, like the root of trust is going to differ very much based on just where the person is living. If that's the, if the root of trust needs to be a government issued document. And maybe again, maybe you don't want to have those restrictions on that either.

Eric Scouten (37:19.17)
I don't think that we can practically because, you know, many people live in jurisdictions that aren't issuing digital credentials at this time. So you know, a necessary part of it from our perspective is working with KYC vendors who can provide an alternative that takes your, you know, traditional non-digital identifier and turns it into something that they can attest to in the digital world.

Mathieu (37:44.15)
And in the spirit of sovereignty, maybe some folks don't even want to share that or have their content associated with a government identity.

Eric Scouten (37:56.08)
Yeah, yeah. We've spoken before about citizen reporting where that could be potentially hazardous to do. The other example that we think about a lot, and we're still sort of puzzling through how do you express that in this world, is people who intentionally have anonymous or pseudonymous identifiers. The classic example is Banksy. Nobody knows who he really is in real life, so he can't use his government identity for his work, but perhaps he wants to identify as Banksy in some way that is consistent and that he can ensure that nobody else uses that reputation in his name, so forth. But I don't know how you verify a pseudonym. We haven't figured that out yet.

Mathieu (38:46.09)
Yeah, it seems like that last mile problem is a complicated one.

Eric Scouten (38:50.21)
Yeah. So, you know, we're starting from the problems that we know we can solve and trying to make sure that the design decisions that we make in this identity space don't prevent us from solving other problems of that sort at a later time.

Mathieu (39:03.13)
Are there any other participants within the whole content creation ecosystem that we haven't talked about maybe yet, that need to be involved? Do browsers need to be involved? Are there other types of vendors in the space that need to be able to support this thing?

Eric Scouten (39:21.13)
I think the places where we're seeing slower adoption, some of them are camera devices. We may have some hope for that in the near future, but capture devices are still sort of a weak

point. The other place where I guess I would say I'm concerned about it is in social media sites. So a lot of social media likes to strip metadata that they aren't familiar with.

And it would certainly help if those became more proactive about showing people what the history is, if it is available. We've done some prototypes. We have a browser extension that I think runs in the Chrome browser that will show you that provenance sort of outside of that, but it's still something you have to go out and explicitly install. It isn't natively part of the browser at this time.

Eric Scouten (40:14.13)
So we hope that people's interest in this space and in understanding how the content gets to them will drive increasing adoption in these parts of the overall ecosystem.

Mathieu (40:28.03)
Yeah, it sounds like any social media platform or distributor that's focused on the truth. And I think about the, I guess the mission or statements made by Elon Musk and X would be, if those values align quite nicely to this, it's like, why wouldn't they want to do this? It's just, again, for them would increase the value of their network.

Mathieu (40:50.19)
are there any other topics that I want to, I don't know if it makes sense to go more technical and anything else or anything we haven't covered, Eric, that you think would be worth covering? I'm just looking through my notes.

Eric Scouten (41:04.01)
Yeah, I think I would extend an invitation to people to study the work, particularly that I'm leading in the Creator Assertions Working Group. If they're interested in individual and organizational identity, that's a fairly easy group to join and would be interested in experts in the field being part of that process. And I can send you links for how to do so.

Mathieu (41:32.01)
Yeah, we'll post those. It seems like there's no new like, there's no new groundbreaking technology here, right? That's just assembling existing pieces into a new architecture type of thing. And obviously getting the buy-in from the right stakeholders, but is that a fair statement?

Eric Scouten (41:49.23)
That's actually very well aligned with some of the decisions that we made early on in the architectural process is to try not to invent new things as much as possible, but to use existing and understood technologies and put them together in new ways.

Mathieu (42:04.02)
Are there case studies or things that have gone live so far? And then maybe an extension on that question is, is there type of content that you're expecting to see in the near future that will have this stuff that will start to showcase how it could solve some problems?

Eric Scouten (42:22.18)
Well, it's very interesting you mentioned that literally earlier today as we're recording I got word from the BBC that they are now publishing content credentials on their website. So that's a very big Demonstration of adoption. I believe there's a Chilean newspaper that did something similar a couple of months ago. So the hope, I think, is that more and more news media pick up on this and use that as a way to differentiate their actual reporting from that which would misrepresent their messages.

Mathieu (43:00.17)
Congratulations on that. I'm looking forward to seeing that. So as a consumer, I'd be able to see in the browser the icon on content, and we'd be able to see a provenance of that.

In terms of visualization of this, you mentioned a tool earlier, it may be worth us posting this as well. So is there like a unique identifier that's generated by the standard for every piece of content that could then be dropped in and queried? And I guess it also depends where things are being stored.

Eric Scouten (43:35.13)
Right, so in general, the C2PA manifest store that I talked about, the collection of manifests that describes the current asset and any prior assets that were incorporated into it, can either be embedded directly in a file, or it can be referenced from the file and stored in some kind of cloud storage. That's basically an HTTP hash link. Um, there is not a central, uh, registry of content. So, um, you know, the, the kinds, you know, a query of does something with this, you know, global unique identifier or whatever exists is not feasible. That said, we have done at Adobe, uh, project where content that is signed through our service can be stored on what we call our manifest cloud.

Eric Scouten (44:30.11)
which is basically, you know, any Adobe authored content can have its manifest detached and placed online, and we do have some options there for visual similarity searches that would allow you to say, I have this piece of content, I don't know where it came from. Can you find something similar with the C2PA manifest and use that to describe what I'm looking at? And there's work being done in the C2PA to see if that approach can be converted from something that's currently proprietary at Adobe to something that's more accessible on generic standards.

Mathieu (45:06.03)
I guess that distributed architecture works quite well with how the internet works too. And this is where I'm excited about where trust registries come into this as well, because when you're looking for more context around a specific domain type of thing, it's quite impossible for any particular party to provide that service. It's kind of an all in one. So just like you're getting an input here on the provenance of the content, you may want to get inputs on, you know.

the signatures behind it and are these actually, according to an authority under control of a specific organization or person. And I actually personally got into the whole trust registry space from the angle of digital credentials, but the more and more time we spent on just trust lists or governed lists or registries, however we want to call it, starting to see broader applications for it outside of credentials.

Mathieu (46:05.14)
you start to realize it's just another input. It's completely independent of digital credentials and just having trust registries for signatures but like based on or governed under specific domains or specific verticals on the internet are gonna become quite powerful. I think in a world where everything needs to be signed to be trusted, which is, I think where we are today and where we're continuing to go towards. I'm quite excited about the merging of these two things.

Eric Scouten (46:37.20)
Yeah, I'm very thankful to see that interest in trust registry and trust governance come to the forefront in the identity community because it fits very nicely with the needs that we see for understanding who vets the identity of the content creators that you're working with.

Mathieu (46:57.06)
Yeah. And the fact that it's distributed, decentralized, it's kind of like you could use it as an input, but again, it's not, it's not making a decision for you. It's just another data point.

Eric Scouten (47:05.06)
Right. I think one of the challenges that we face in this space is that there are, you know, even in a single manifest, quite a number of data points that are there as trust signals. And how do we communicate those in a way that's approachable and comprehensible to you as a content viewer? It's, you know, if we were to spill the entire report of what's in a C2PA manifest, it would likely be overwhelming, but.. You still need to be able to get to it if somebody has a question of, well, who actually vetted Mathieu's identity? And you want to dig through, you know, whatever the government or KYC vendor, whatever it is that understood your identity and attests to it and who says they're valid and so on and so forth. It's there, but you have to choose carefully how to present that.

Mathieu (47:56.11)
And because it's decentralized distributed, it's not like I'm on a platform and there's a checkmark. And I know there's been a verification, but it's like, it's so contextual. So the user experience definitely is going to be a, an opportunity, let's say.

Mathieu (48:15.21)
Well, thanks a bunch for doing this with me. And it was nice to have you at Excite last week to do this with you today. And then I'm looking forward to this icon presentation that we have on Wednesday.