Contingent Accessibility

Why availability is not access when the interaction itself keeps changing.

THOMAS ROCHA III
MAY 10, 2026

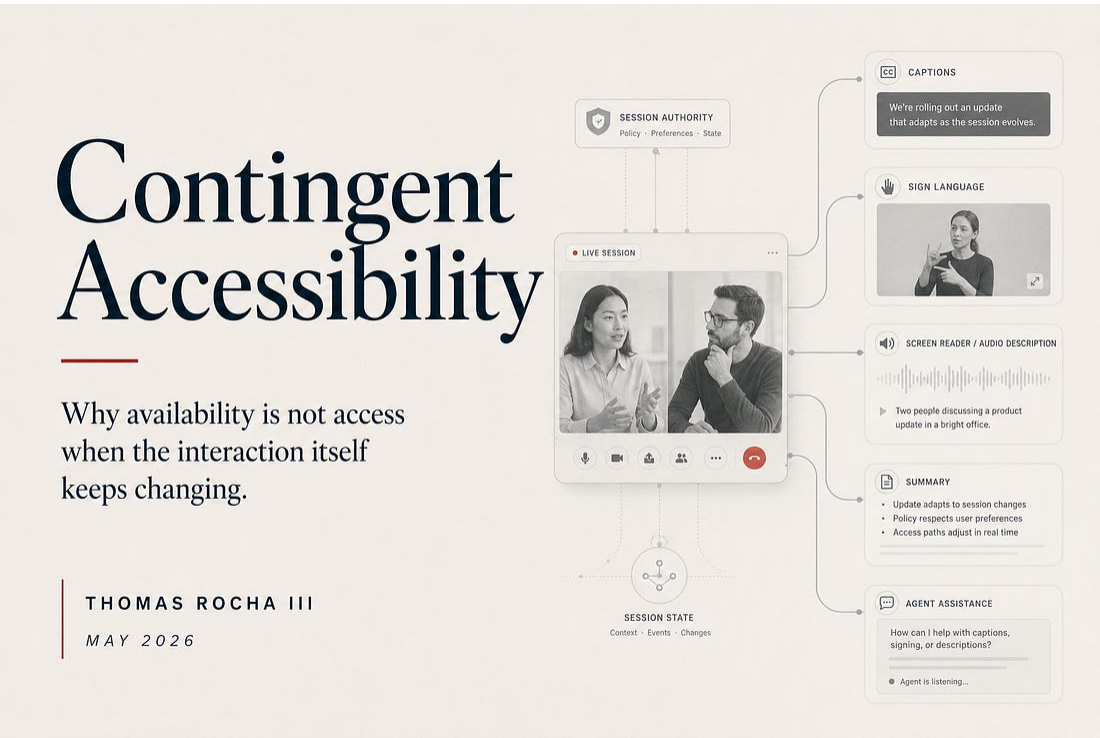

The accessibility front is moving faster than most institutions have noticed, and the strongest signals are not the headline lawsuits. They are the deadline slips, the failed settlements, the shifting litigation patterns upstream, and a regulatory landscape that has begun to attach legal obligations to what happens during execution rather than to what was documented beforehand.

On April 20, 2026, four days before the long-scheduled compliance date, the Department of Justice published an Interim Final Rule (91 Fed.

Reg. 20902) extending the deadlines for ADA Title II web accessibility compliance. State and local government entities serving populations of 50,000 or more were required to bring their web content and mobile applications into conformance with WCAG 2.1 Level AA by April 24, 2026, under the final rule the DOJ issued on April 24, 2024. That deadline is now April 26, 2027. Smaller jurisdictions and special district governments, originally scheduled for April 26, 2027, now have until April 26, 2028. The rule itself did not change. The standard did not change. The obligation did not change. What changed is that the federal government conceded, four days before the date, that the institutions subject to the rule were not ready.

That is not a retreat. It is a failure signal. The DOJ stated plainly that it had overestimated the technological readiness and institutional capacity of covered entities. Two years' notice, fifteen years of prior rulemaking, a published technical standard, and the systems still could not deliver. The deadline moved because compliance is not what most of the field has been doing. Compliance, as the rule actually requires it, attaches to the live experience of a user with a disability interacting with the system. The architecture deployed across most institutions cannot prove that experience.

Meanwhile, the Department of Health and Human Services compliance date for healthcare organizations receiving HHS funding remains May 11, 2026. That deadline has not been extended.

What is worth saying plainly at this point, before the discussion moves to the EU, to California, or to the litigation patterns, is that no one can prove deterministic interaction-level compliance at scale. Not in California. Not in the EU. Not under HHS Section 504. Not under DOJ Title II. Components pass audits. Features ship. Accessibility statements get posted. The thing the rules actually require, which is provable access during the live interaction, is not what any deployed architecture currently produces. The deadline slips are not stories about lagging institutions catching up to leading ones. They are signals that the

leading institutions cannot prove compliance either. Component compliance is what the field calls compliance because it is what the field is capable of. The rules describe something else.

That observation governs both the government posture and the civil one. Federal regulators cannot enforce a standard the architectures cannot meet, and they know it; the DOJ extension is the legible form of that knowledge. Civil plaintiffs cannot recover at the volume the damages structures permit unless the architectures change, and the affirmative-defense bills moving through state legislatures are the legible form of that knowledge from the other side. The current arrangement, in which deadlines slip, settlements are negotiated around best-effort documentation, and litigation focuses on specific barriers rather than systemic enforcement, is a temporary equilibrium. It is stable only while two statements remain true at the same time: no widely deployed system can prove behavior in real time, and no authority can require that it be done. When either statement stops being true, the equilibrium ends.

What the Europeans already enforce

While the American conversation has centered on whether the Title II rule will hold, the European Union has been operating under a more stringent regime since June 28, 2025. The European Accessibility Act, Directive (EU) 2019/882, became applicable across all twenty-seven member states on that date, and it does not stop at public-sector websites. It reaches e-commerce platforms, banking services, transport ticketing, telecommunications terminals, e-books, audiovisual media access, ATMs, point-of-sale terminals, and consumer computing hardware. The technical standard, EN 301 549, incorporates WCAG 2.1 Level AA for web content and adds product and service requirements that go well beyond it.

A United States company that accepts orders from a customer in Lyon or Lisbon falls under EAA jurisdiction. The directive covers

manufacturers, service providers, importers, and distributors, with liability that can attach at multiple points in the supply chain. Each member state sets its own penalties, which must be effective, proportionate, and dissuasive. Enforcement authority began when the rule became applicable, and member states now have the legal machinery to treat accessibility as a market-access condition.

The EAA is the more architecturally important regime because it is cross-domain. A vendor can pass a website audit and still fail an EAA compliance review, because the obligation reaches the experience across components: the booking flow, the payment terminal, the consumer device, the audio-visual stream, the customer service channel. Component compliance is not the same as service compliance. The EAA was the first major regulatory regime to operationalize that distinction.

Litigation pressure is moving upstream

In California, the Unruh Civil Rights Act provides what the federal ADA does not: a private right of action with statutory damages. A plaintiff can recover a minimum of four thousand dollars per violation, plus attorney's fees, with treble damages available where intent can be shown. That damages structure is the engine of the country's largest concentration of accessibility litigation.

The legislature has been trying to recalibrate the litigation economics for several sessions. Assembly Bill 2190, currently moving through the 2025-2026 session, would create an affirmative defense to statutory damages claims if a business can show, within thirty days of a pre-lawsuit demand, either that it had published a digital accessibility report disclosing the specific barrier and updating it with remediation, or that it had conducted regular automated and manual testing in good faith. The defense is conditioned on actual evidence of behavior, not on the existence of a policy document.

The bill is contested, including by accessibility advocates who argue the affirmative defenses set traps for businesses that disclose their issues and reward those that stay silent. That debate matters less, for the architectural argument, than what the bill assumes about evidence. AB 2190 takes for granted that the relevant proof is dynamic and ongoing: monitoring frequency, remediation timelines, third-party component governance, public disclosure that updates as conditions change. That is the compliance posture the regulators want and the architectures cannot supply.

The bill also reaches resource service providers. In its predecessor (AB 1757) and in the current draft, the legislature is moving liability upstream toward the developers, vendors, and component providers whose products are embedded in the noncompliant experience. California is no longer treating accessibility as the sole responsibility of the merchant. The component layer is in scope. The integration layer is in scope. The session in which the access either occurred or did not occur is in scope.

The Fashion Nova irony

On February 2, 2026, the Department of Justice filed a Statement of Interest in a federal court in Oakland opposing a proposed five point one five million dollar class settlement of a website accessibility case against Fashion Nova. The DOJ's objection was not that the settlement was too small. It was that the settlement website itself was allegedly inaccessible to the class members it was supposed to compensate. The agreement also lacked durable forward-looking enforcement, and the attorneys' fees exceeded the residual recovery available to the class.

That filing is the cleanest single illustration of what is failing. A system designed to remedy an accessibility failure reproduced the failure on the remediation surface. Not because anyone intended the irony. Because the architecture used to build the remediation site was the same architecture that produced the original violation: feature

compliance without interaction governance, component testing without session evidence, a static accessibility statement standing in for live, provable access.

That pattern is not unique to Fashion Nova. It is common, and it is the deeper reason the deadlines slipped. Institutions cannot meet the compliance obligation because the architectures they have deployed cannot generate the evidence the obligation requires.

Availability is not access

The first mistake is thinking accessibility means adding something. Captions. A transcript. A screen reader label. A signing window. A larger font. A contrast setting. An AI summary. Each of these is real. Each may be necessary. None of them is the architecture.

The architecture is what decides when the accommodation is required, who is entitled to receive it, whether it is available, how it is delivered, what happens when the session changes, and how the system proves the participant received access comparable to everyone else. That is the part most systems do not have. They have features. They do not have accessibility authority.

A platform can generate captions from shared meeting audio. A participant can turn them on. The feature appears individualized because the display is individualized. But the underlying architecture is usually not individualized in the meaningful sense. The session has a caption artifact. Participants may expose it. That is useful, but it is not the same as governing an accommodation.

A caption toggle does not answer the harder questions. Did the participant need captions? Did the participant receive them continuously? Were they accurate enough for the context? Did they persist through reconnect, handoff, escalation, device change, or breakout? Were they available without forcing disability disclosure?

Were they logged as an accommodation rather than as a broadcast attribute? Could the participant have received something better if something better had been available?

The platform can say captions exist. The law increasingly asks whether access existed. Those are not the same claim, and the difference is what the regulatory architecture is now trying to evaluate.

Special needs became ordinary features

There is a quieter failure happening alongside the loud one. The system exposes an auxiliary layer to everyone and calls it accessibility.

Everyone can turn on captions. Everyone can turn on summaries. Everyone may soon be able to turn on translation, signing avatars, visual description, coaching, sentiment analysis, or agentic note-taking. The feature is available to all. The accommodation becomes indistinguishable from preference.

That is convenient for platforms because preferences are local. A user setting can be turned on or off without changing the authority model of the session. But an accommodation is not a preference. It is an entitlement attached to a participant in a live interaction. It may be legally required, medically necessary, educationally mandated, or a condition of employment, civic participation, court access, healthcare access, or public service access. Treating that as a display preference collapses the category.

The problem is not that everyone may use the feature. The problem is that the system no longer knows when the feature is convenience and when it is required access. That distinction matters because convenience can fail gracefully. Access cannot.

Deaf access exposes the media problem

Deaf and hard-of-hearing access makes the issue legible, because the auxiliary stream is obvious. Captions are text derived from audio. Signing may be human video, pre-recorded library content, stitched signing libraries, or AI-generated avatar output. Translation may require speech recognition, semantic conversion, sign-language grammar, regional dialect, facial expression, motion rendering, and timing.

The moment signing enters the session, the caption-as-feature model begins to break. A platform cannot generate every sign-language stream in every language for every participant just in case someone needs it. It cannot send every interpreter feed, every avatar, every translation, and every fallback to every endpoint as a universal broadcast.

Availability cannot mean transmission. The law may push toward universal availability. Physics forbids universal broadcast.

The only workable structure is selective delivery under session authority. A participant signals need. The system determines entitlement. The system determines what is available. The system selects the best modality the moment can support. The system routes the stream only to where it belongs. The system preserves it through the session's mutations. The system tears it down when it is no longer needed. The system proves what happened.

That is not a media feature. It is a governed state transition.

Vision access exposes the state problem

Where deaf access exposes auxiliary media, blind and low-vision access exposes something deeper. It exposes session-state failure.

A blind participant does not only need words spoken aloud. They may need to know who joined, who left, who is currently speaking, who raised a hand, who shared a screen, what slide is visible, what chart changed, what button became active, what poll opened, what chat

message arrived, what annotation was drawn, and whether the host changed permissions.

That is not content. That is the live state of the interaction.

A screen reader can read a button if the button is labeled. It cannot guarantee access to a live meeting if the meeting state is not semantically governed. A visual-description agent may describe a screen share. But who authorized that agent to inspect the shared content? Where is the description processed? Who receives it? Is it retained? Does it persist when the presenter changes windows? Does it stop when the participant leaves the room? Does it follow the participant from laptop to phone, with the same authority?

The thing that has to be made accessible is not a page. It is the changing state of the interaction itself. A static accessibility layer cannot govern that.

The agent problem makes everyone special

Accessibility is the first visible case of this architecture problem. Agents are about to make it universal.

A platform summary is one auxiliary artifact. A meeting transcript is another. A participant-side note-taker is another. A sales coach, a legal assistant, a medical documentation tool, a personal memory agent, a translation agent, an accommodation agent. A ten-person meeting with two agents per participant is no longer one meeting. It is one interaction spawning twenty auxiliary participants, twenty derivative records, twenty memory surfaces, twenty possible action trails, and twenty separate authority questions.

Who captured the session? Who summarized it? For whom? Stored where? Under whose policy? Retained how long? Shared with whom? Allowed to act downstream? Subject to privilege, to disability

accommodation rules, to medical or educational or employment record obligations?

A recording banner does not answer any of this. A consent checkbox does not answer it. A meeting setting does not answer it. The system has moved beyond recording into producing derivative state, and the derivative state outlives the call.

Once every participant can bring an agent, every participant can create a private fork of the interaction. Some forks are harmless. Some are regulated. Some are privileged. Some are accommodations. Some are surveillance. Some are memory. Some are authority-bearing action that affects downstream systems. The platform cannot govern that by pretending the meeting ended when the call ended. The interaction continues through its artifacts, and so does the obligation.

The compliance surface has moved

The old compliance question was simple. Does the product support accessibility?

The new question is harder. Did the live interaction provide effective access to the participant who needed it, using the best available accommodation, while preserving authority, privacy, continuity, and proof?

That question cannot be answered by a feature list. It cannot be answered by a VPAT. It cannot be answered by a post-hoc log assembled from disconnected systems after the session ended. It has to be answered during execution, in the same window the action occurred, by something that knew at the time of the decision what was admissible.

The DOJ extension is, in part, a recognition of this. The agency conceded that institutions could not meet the deadline because the underlying architectures cannot produce the proof the rule

contemplates. The EAA is operationally already there: enforcement actions across the EU since June 2025 have surfaced cases where vendors had passed component audits and still produced an inaccessible service experience, and member state regulators are following the experience, not the audit. California AB 2190 takes the same posture, requiring evidence of behavior over time rather than evidence of a policy document. The Fashion Nova filing is the case where even the remediation surface failed the test the underlying suit was about.

The burden is shifting from documentation of controls to proof of behavior. That shift exposes the architectural weakness directly. If a system cannot distinguish vendor capability from institutional authority during the live interaction, it cannot prove that its own policy governed the access. It can only prove that a feature existed somewhere upstream.

The missing layer

The missing layer is not another caption engine, another interpreter marketplace, another summary tool, another overlay, or another audit dashboard. It is the authority that binds accommodation, auxiliary computation, participant entitlement, and derivative artifacts to the live interaction itself.

That authority must sit outside the application feature and outside the transport path. The transport carries the signal. The application exposes the function. The session governs whether the function is admissible for this participant, in this context, under this policy, at this moment. Without that boundary, every accommodation becomes a feature, every feature becomes a stream, every stream becomes an artifact, every artifact becomes a policy problem, and every policy problem becomes after-the-fact reconstruction.

That is why the current platforms cannot honestly claim deterministic compliance with what the regulations now require. They can comply at the feature layer. They cannot prove what happened at the interaction layer.

Accessibility is the proof case

Accessibility is not the edge case. It is the proof case. It shows what happens when a live interaction has to deliver different experiences to different participants under different obligations without breaking the common session.

The hearing participant receives audio. The Deaf participant may require captions or signing. The blind participant may require semantic state narration. The mobility-impaired participant may require keyboard control or voice navigation. The neurodivergent participant may require pacing, summarization, or reduced sensory load. The participant using an agent may require translation, memory, action extraction, or regulated recordkeeping.

Each may be entitled to a different auxiliary path. Each path must remain bound to the same interaction. Universal availability. Selective delivery. Continuous authority. Provable execution. That is the only structure that scales.

Accessibility does not ask whether a system has features. It asks whether a person could participate. In modern systems, participation is no longer a property of the interface alone. It is a property of the session.

A session without authority cannot prove access.

A final observation is worth making, because it changes how the regulatory landscape should be read.

The DOJ extension, the EAA posture, AB 2190, the Fashion Nova objection, and the broader settlement pattern all operate inside the same assumption: no widely deployed architecture can yet prove access during execution.

That assumption is doing the work.

It is why deadlines slip. It is why best-effort documentation survives. It is why regulators describe the gap as readiness rather than architecture.

The moment a viable architecture exists, the posture changes.

Best effort narrows.

Documentation stops being proof.

Feature availability stops being access.

Compliance becomes a measurable state.

That shift has not happened yet.

It is the shift this architecture is built to force