NOVARA
CONSULTING GROUP

A Court Just Wrote Rules for Machine Captions. Nobody Has Written Them for Machine Signing.

Heather Grizzle · November 18, 2026

The SiriusXM decree makes every AI transcript say how it was made and lets Deaf plaintiffs pick what gets human review. Machine signing has no such rules.

Contents

1. Five years to get a transcript
2. What the decree requires
3. The clause that matters most
4. Where the decree is thin
5. Machine signing has no decree
6. What any buyer can borrow tomorrow
7. The benchmark
8. Sources

On 16 November 2026, after five years in federal court, SiriusXM and Pandora signed a consent decree with the National Association of the Deaf over podcast transcripts. Strip away the podcast details and what remains is rare: a set of court-enforceable rules for an AI system that serves Deaf people. Say how each transcript was made. Name the benchmark. Give users a button to report errors. Let the Deaf plaintiffs choose which shows get a human check. Machine-generated sign language, which is being installed in train stations, airports and websites right now, has none of this.

IN BRIEF

- The consent decree in National Association of the Deaf v. SiriusXM was filed in the Southern District of New York on 16 November 2026 and is awaiting the judge's approval. It covers transcripts for English-language podcasts on SiriusXM, Pandora and Stitcher for 36 months.
- Every transcript must state on its face how it was made, for example by automatic speech recognition alone or with human edits. That single clause is the most important precedent in the decree for anyone governing accessibility AI.
- The decree pegs accuracy to a named AI model, ElevenLabs Scribe v2, as of the day the decree takes effect. It sets no number of its own.
- Human review applies to 30 podcasts a year, chosen by the Deaf plaintiffs from the 50 that earn the most money. Everything else is machine output, labeled as such.
- Quality reports go to the plaintiffs' lawyers, not the public, and the company can propose ending human review when it judges the AI good enough. Those are real limits. They are still more governance than any sign language AI product I have reviewed this year is under.

THE DECREE, BY THE NUMBERS

5 years
from the December 2021 complaint to a signed decree

36 months

30 of 50
top-earning podcasts that get human review



term of the decree once the judge approves it

each year, picked by the Deaf plaintiffs

1
named AI model that sets the accuracy standard

$1.465M
attorneys' fees and costs, against $35,000 in damages

Source: Consent Decree, National Association of the Deaf v. SiriusXM Holdings Inc., No. 1:21-cv-10542-JAV (S.D.N.Y., filed 16 November 2026)

Five years to get a transcript

The podcast boom happened without Deaf people in it. For most of the past decade the biggest audio platforms in the United States offered millions of hours of talk, interviews, news and comedy, and no way to read any of it. In December 2021 the National Association of the Deaf, represented by Disability Rights Advocates, sued SiriusXM, Stitcher and Pandora under the Americans with Disabilities Act and New York state and city human rights law. The complaint's point was simple. The companies were selling a service to hearing people that Deaf and hard-of-hearing people could not use, and transcripts would fix it.

The case took five years. It changed judges. The companies moved for judgment on the pleadings. The parties went through three rounds of private mediation before a retired judge, Steven M. Gold. In October 2026 the plaintiffs' lawyers told the court an agreement was close, and on 16 November the signed decree was filed for Judge Jeannette Vargas to approve. The four individual plaintiffs, Jazmine Jones, Amber Martin,

James Munro and Mei Nishimoto, signed it in the first week of November. James Munro dates his part of the fight to June 2021.



Figure 1. Five years of litigation, then a three-year decree. The clock on every obligation starts when the judge signs.
Novara Consulting Group, from the court docket and the consent decree

I want to be clear about why I, a DeafBlind woman who works on sign language AI, am writing about podcast transcripts. It is not because transcripts and signing are the same thing. They are not, and a transcript in English is no substitute for a Deaf person's access in ASL. It is because this decree is a working model. It shows, in fourteen pages of enforceable text, what it looks like when Deaf people get to set terms for an AI system that is supposed to serve them. That model is badly needed one aisle over.

What the decree requires

The core of the deal is a two-tier transcription system. For podcasts SiriusXM owns or sells advertising on, the company must transcribe every new English-language episode using ElevenLabs Scribe v2, an automatic speech recognition model, and then run the output through a Transcript Enhancement Service it must use reasonable best efforts to build within six months. That service adds what raw speech recognition misses: speakers labeled by name rather than as Speaker 1 and Speaker 2, sound effects and laughter and music descriptions, which the decree calls atmospherics, and proper punctuation, capitalization and paragraphs. For third-party podcasts the company only licenses, transcription goes through Rev.

Around that machine core, the decree builds a set of obligations that any procurement officer will recognize as governance. The company must write a style guide covering speaker labels, sound descriptions, filler words, stuttering, censored words and long single-speaker content, and send it to the plaintiffs within fourteen days of the decree taking effect. The Deaf plaintiffs get thirty days to comment. Where the company rejects a suggestion, it must explain why in writing.

A person must review and edit the machine transcript for thirty podcasts a year. The company hands the plaintiffs its list of the fifty English-language podcasts that earn it the most revenue, and the plaintiffs pick the thirty.

Every transcript page must carry a link reading "Report a problem with the transcript." The company must log every complaint, say how it resolved each one, and summarize them in an annual report.

Once a year the company must evaluate a sample of transcripts against named metrics: word error rate, grammar, accuracy of speaker labels, accuracy and placement of sound descriptions, and conformance with the style guide. It must survey the market for the best available transcription vendor. It must use custom word lists for each podcast to improve accuracy.

Staff who build the system or answer transcript complaints must take annual training on Deaf culture and communication access, delivered at least in part by a Deaf organization or a Deaf instructor. And the plaintiffs get guest accounts, so they can check the product themselves.



The clause that matters most

One paragraph stands above the rest. Paragraph 27 requires the companies to ensure that every transcript "identify on the Transcript itself the manner in which the Transcript was created," and gives the examples of speech recognition with human edits or speech recognition alone. The company chooses the wording. It does not get to choose silence.

That is a disclosure rule for AI output, written into an enforceable court order, for the people who depend on that output. A Deaf reader looking at a transcript will know whether a person checked it. That changes how much weight they put on it. A transcript of a medical podcast that says it was machine-made alone is a transcript to read with care.

Now look at sign language AI. A Deaf passenger watching a signing figure on a train platform, a Deaf patient watching an avatar on a hospital website, a Deaf student watching a signing plugin on a course page: in almost every deployment I have examined this year, nothing on the screen tells them whether a person who knows the language reviewed that signing before it reached them, whether it was recorded from a human interpreter, or whether it was generated by a model on the fly. In my piece on Deaf AI's airport pitch, the safeguard the company described to a student newspaper was the one it offered airports as optional. Paragraph 27 is what it looks like when that choice is taken away from the vendor.

Where the decree is thin

None of this means the decree is a model to copy word for word. It is a negotiated settlement between a large company and a small coalition, and it shows. The limits are worth naming, because they are the same places a sign language AI contract will be pushed to give way.

The accuracy standard is a brand name, not a number. Paragraph 30 sets the bar for owned and licensed podcasts as an accuracy rate equal to or greater than ElevenLabs Scribe v2 "as of the Effective Date." The decree never says what that rate is. On an independent benchmark published in March 2026, Scribe v2 ranked first with a word error rate of 2.3 percent, which is genuinely good. But a podcast is not a benchmark, and a standard defined as whatever one vendor's model did on one day is hard to test and easy to argue about. For third-party podcasts the bar is whatever Rev offered on



that day, and the decree does not say which Rev product it means. Rev sells both human and machine transcription. Its machine service had the highest error rate shown in that same ranking, 6.1 percent.

Human review follows the money. Thirty shows out of the fifty that earn the most is a sensible way to put people where the audience is. It also means the smaller shows, which are often where Deaf, disabled and minority creators are, get machine output only. And the plaintiffs, not the company, have to make the hard choice of which twenty of the top fifty get no human check at all.

The company can propose ending human review. Under paragraph 33(c), if SiriusXM concludes from its own evaluations that the AI has improved enough, it can propose sunsetting human cleanup. The plaintiffs can object, but they must explain in writing within thirty days how continued review would materially improve quality. They get no discovery of the company's internal data to do it. The party that saves money by ending review holds the evidence, and the burden falls on the Deaf plaintiffs to rebut it.

The results stay private. The annual evaluation and the complaint summaries go to the plaintiffs' lawyers. The public, including every other podcast platform the NAD hopes will treat this as a benchmark, will not see the word error rates or the complaint counts.

Coverage depends on consent. The decree defines a podcast as content for which the companies have "the express consent of the content creators to create transcriptions." A creator who declines falls outside it.

The style guide does not bind. Paragraph 29 says the style guide creates no legal obligation beyond what the decree itself says, and the company keeps final say over its wording.

And it ends. The term is thirty-six months, extended only if the build of the transcript system runs late.

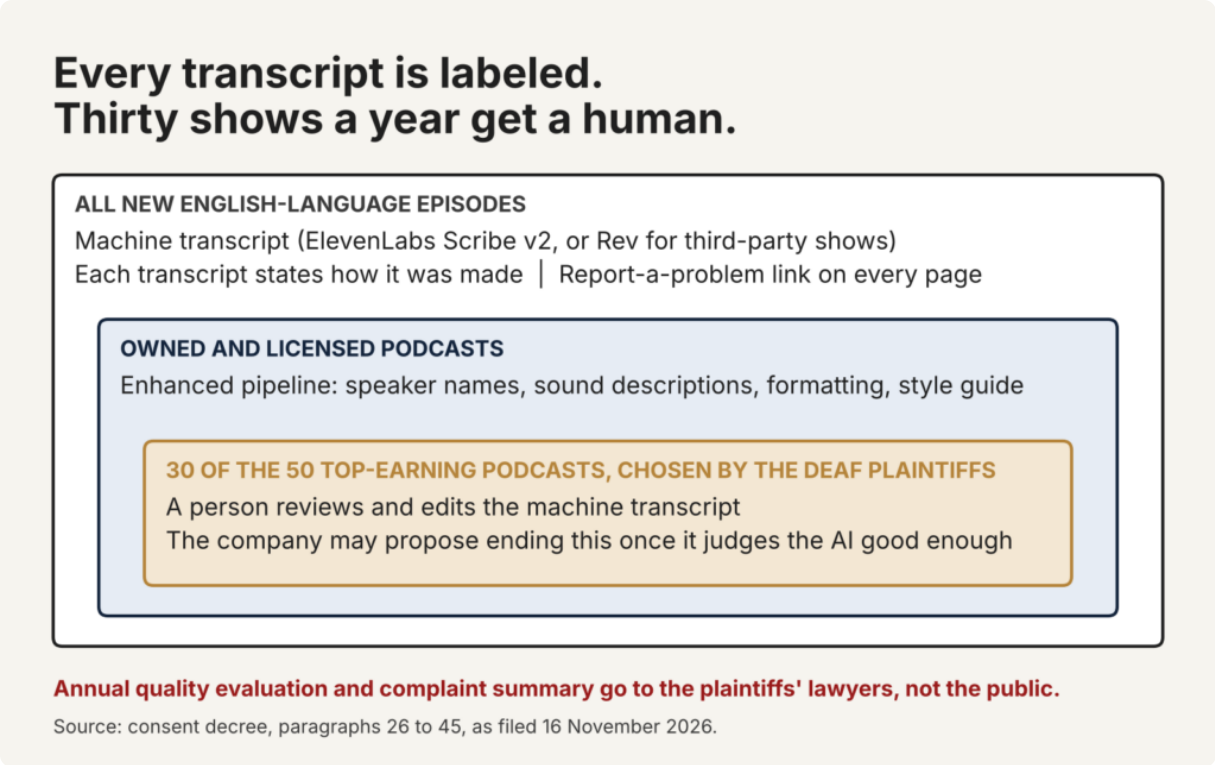


Figure 2. Three layers of coverage. Every transcript is labeled, only some are enhanced, and thirty shows a year get a human.
Novara Consulting Group, from the consent decree

I list these not to knock the people who won this. Five years of litigation for $35,000 in total damages and a three-year order is a hard road, and the plaintiffs and their lawyers got things into this decree that no regulator has required. I list them because each limit is a pressure point, and anyone writing the first equivalent for sign language AI should see them coming.

Machine signing has no decree

Put the decree's main provisions next to the sign language AI market and the gap is plain. The table below sets them side by side, provision by provision.

What the SiriusXM decree requires of machine captions, and what exists for machine signing

Decree provision	What it requires for transcripts	Equivalent for sign language AI today
Disclosure of method (para. 27)	Every transcript says how it was made	No general requirement; rarely shown on screen in deployments NCG has reviewed
Named accuracy benchmark (para. 30)	At least as accurate as a named model on a fixed date	No agreed measure of signing accuracy; no vendor NCG reviewed has published an independent evaluation of its signed output
Human review chosen by Deaf people (para. 33)	Deaf plaintiffs pick 30 shows a year for human editing	Human review offered at the vendor's or buyer's option where offered at all
Report-a-problem link (para. 38)	On every transcript page; complaints logged and resolved	Rarely present; Deaf users usually cannot report a wrong sign from the screen
Annual evaluation with named metrics (para. 44)	Word error rate; speaker labels; sound descriptions; style	No standard metric set; output validation rarely specified in contracts
Deaf-led training (para. 37)	Delivered at least in part by a Deaf organization or instructor	Not a standard contract term
Style guide with Deaf input (para. 23)	Plaintiffs comment; rejections explained in writing	No equivalent for how signed output should look or what it must carry
Enforcement (paras. 47 to 50)	Court keeps jurisdiction to resolve disputes	Depends on a buyer writing it into a contract or on a complaint after harm

Source: Consent decree, NAD v. SiriusXM (filed 16 November 2026); NCG public-evidence reviews of sign language AI vendors, 2026

The middle column took five years of litigation to get. The right-hand column is the state of a market that is already selling into public transport, government and education. When I filed an ADA complaint over AI-generated ASL on the train screens at Cincinnati's airport in October, there was no label telling a Deaf traveler how that signing was made, no button to report it, and no evaluation anyone could point to. The airport took it down fourteen days later, but only because a complaint forced the question.

The technical gap matters too. Speech recognition has word error rate, a crude but agreed measure that a court can write into an order. Sign language AI has no equivalent that Deaf people, vendors and buyers have accepted. That is not an excuse for doing nothing. It is the reason the people who use the language have to be the ones who judge it, which is the principle at the center of the SLAT work I publish. The SiriusXM decree reaches for that principle with its style guide comments, its plaintiff-chosen review list and its guest accounts. A sign language version would need it more, not less.

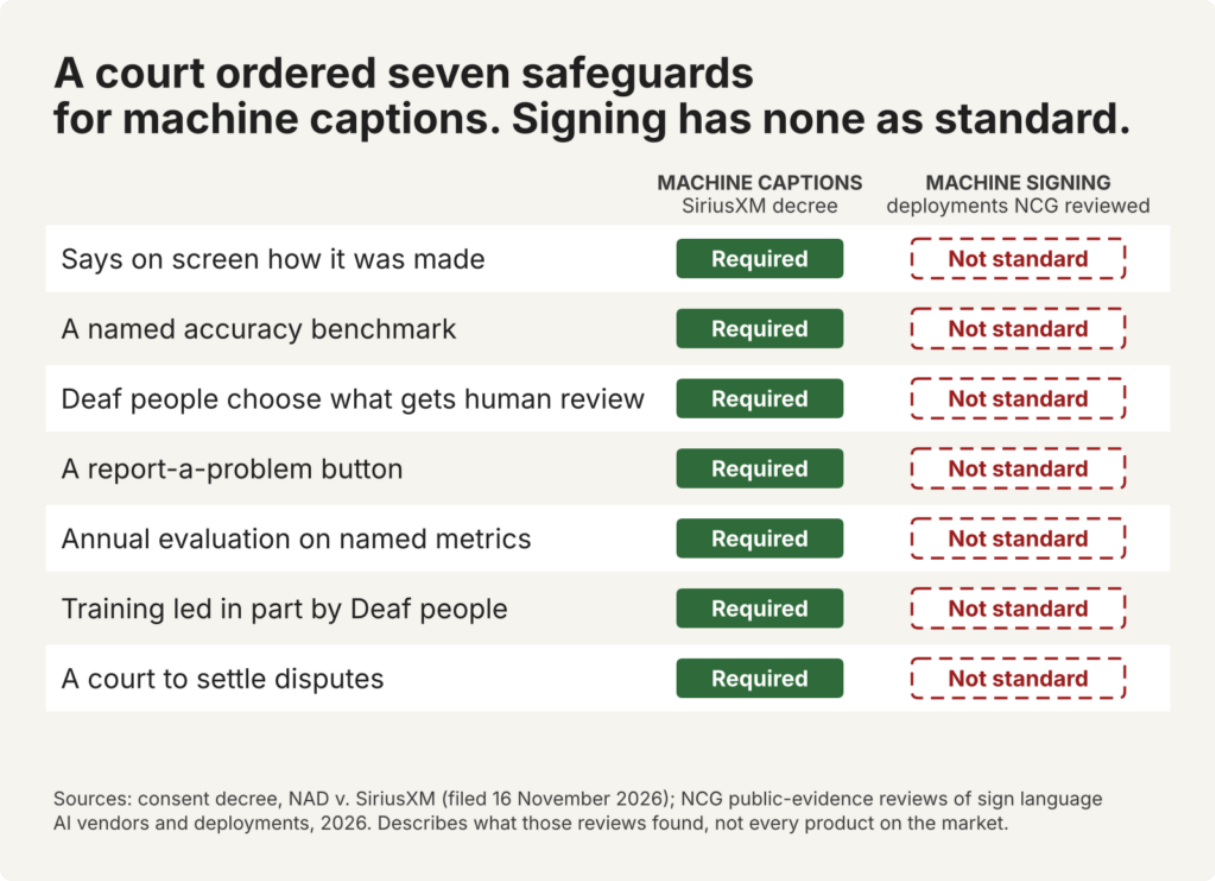


Figure 3. Seven safeguards a court just ordered for machine captions. In the sign language AI deployments NCG has reviewed in 2026, none is standard.
Novara Consulting Group analysis

What any buyer can borrow tomorrow

A transit agency, a hospital system, a state agency or a university does not need a lawsuit to put these terms into a sign language AI contract. The decree is a public document. Its structure can be lifted almost directly.

1. Every piece of signed output MUST state on screen how it was made: recorded from a human signer, assembled from recorded clips, or generated by AI, and whether a fluent signer reviewed it before release.
2. The contract MUST name an accuracy standard that can be tested, and MUST say who tests it. Where no agreed metric exists for signing, comprehension testing with Deaf users fluent in the target language MUST stand in for one.

3. A Deaf-led body, not the vendor, MUST choose which content gets human review, and the vendor MUST NOT be able to end review on its own evaluation alone.
4. Every screen or page with signed output MUST carry a way for a Deaf user to report a wrong or unclear sign, in the user's language, with complaints logged and resolutions reported.
5. The vendor MUST deliver an annual evaluation against metrics agreed in advance, and the buyer SHOULD publish the results.
6. Staff who build or support the system MUST receive training delivered at least in part by a Deaf organization or Deaf instructor.
7. The contract MUST give a Deaf-led group standing to raise disputes and a route to resolve them, rather than leaving enforcement to a complaint after harm.
8. A style guide for signed output, covering which language variety is used, how names and numbers are rendered and what the system does when it cannot sign something, SHOULD be written with Deaf input, with any rejected input explained in writing.

Each of those terms has a precedent now. A buyer who is told they are unworkable can point to a public company that agreed to their caption equivalents in federal court.

The benchmark

Nick Kiego, the NAD's chief executive, said he hopes the settlement becomes "a benchmark for other providers of podcast entertainment." It should be read more widely than that. The decree is the clearest public statement I have seen of what accountability for an accessibility AI system can look like when Deaf people are at the table with lawyers and a judge behind them: say how it was made, measure it, let us pick what gets checked, let us report what is wrong, and let a court settle the argument.

Podcasts got there through five years of litigation. Signing does not have to wait that long. Every one of those terms can go into a contract on the next purchase order, and every buyer of sign language AI should ask why they are not already there.

METHOD AND DISCLOSURE

Novara Consulting Group is independent of SiriusXM, Pandora, Stitcher, ElevenLabs, Rev, the National Association of the Deaf, Disability Rights Advocates and every sign language AI vendor, and holds no financial interest in any of them. This piece is based on the consent decree as filed on 16 November 2026, which awaits the court's approval, together with public reporting and press materials checked on 18 November 2026. Statements about the sign language AI market draw on NCG's public-evidence reviews of vendors during 2026 and describe what those reviews found, not every product on the market. Nothing here is legal advice.

Sources

All sources were checked on 18 November 2026.

United States District Court for the Southern District of New York. (2026, November 16). Consent Decree, National Association of the Deaf et al. v. SiriusXM Holdings Inc. et al., No. 1:21-cv-10542-JAV. Hosted by Disability Rights Advocates.

Disability Rights Advocates. (n.d.). National Association of the Deaf v. Sirius XM. Case page.

Disability Rights Advocates. (2026, October 23). Letter to Judge Jeannette A. Vargas re consent decree timing. Via Podnews.

National Association of the Deaf and Disability Rights Advocates. (2026, November 18). SiriusXM and Pandora agree to provide high-quality podcast transcripts for Deaf and hard-of-hearing consumers. Press release via Podnews.

Bonk, L. (2026, November 18). SiriusXM and Pandora will provide podcast transcripts for the deaf and hard-of-hearing community. Engadget.

Radio Ink. (2026, November 18). SiriusXM, Pandora settle deaf podcast transcript lawsuit. Radio Ink.

Podcast News Daily. (2026, November 18). SiriusXM agrees to $1.5 million settlement in podcast transcript fight. Podcast News Daily.

The Decoder. (2026, March 1). ElevenLabs and Google dominate Artificial Analysis' updated speech-to-text benchmark. The Decoder.

Federal Communications Commission. (2014, February 24). Closed Captioning Quality Report and Order, FCC 14-12. FCC.

Novara Consulting Group. (2026, October 28). Human Review, Optional: Deaf AI and the Airport Sign Language Pitch. NCG Insights.

Novara Consulting Group. (2026, October 26). I Filed One Complaint. An Airport Took Its AI Signer Down.. NCG Insights.

Novara Consulting Group. (2026, September 27). The Corpus Question. NCG Insights.

TIPS

Know something about a sign language AI system that the public should know? Send it to tips@novaracg.com. Every tip is read by me and kept confidential. If you work in the UK, writing to us is not by itself a protected disclosure under the Public Interest Disclosure Act 1998, so before you share employer information, check your options with Protect (protect-advice.org.uk), the free whistleblowing advice line.

SUPPORT NCG

NCG does this work without vendor money, so our findings stay ours. If it has been useful to you, you can support NCG's independent work here. NCG is a for-profit company, not a charity. Contributions are not tax-deductible and buy no service or influence.

How to cite

Grizzle, H. M. (2026, November 18). A Court Just Wrote Rules for Machine Captions. Nobody Has Written Them for Machine Signing.. Novara Consulting Group. https://www.novaracg.com/2026/11/18/a-court-just-wrote-rules-for-machine-captions-nobody-has-written-them-for-machine-signing/