NOVARA
CONSULTING GROUP

Compared to Nothing: The Baseline Swap in Sign Language AI Sales

Heather Grizzle · November 1, 2026

A Signapse presenter asked buyers to judge AI signing against no access. The law already set the baseline, and CVG shows why wrong signing is not zero.

Contents

1. Two sources, one platform
2. A request to lower the bar
3. What a baseline swap does
4. The law has already chosen the baseline
5. The law has done this before
6. Wrong signing is not zero
7. The community answered the comparator question in 2018
8. What the vendor's own table shows
9. The record, in order
10. Who holds the duty when the vendor is wrong
11. The quiet part
12. What a fair comparison would look like
13. Writing the comparator into the evaluation
14. Something is not better than nothing
15. References

Two sources, one platform

Picture the platform before picturing the procurement. A Deaf traveller steps off a moving walkway onto a terminal train platform at Cincinnati/Northern Kentucky International Airport. She has done this many times. The train runs between the parts of the airport she needs to reach, and the platform screens tell riders when the next car is coming, where it is going, and what to do if something goes wrong. For most pas-

sengers those screens are a back-up to the announcements overhead. For her they are the whole of the information, because the announcements overhead do not exist for her at all.

On these screens there are two versions of the same message. One is English text. The other is a signed video, generated by software, meant to carry the same message in American Sign Language. The arrangement is meant to be reassuring: the airport has gone beyond text, and the signing is there so that a passenger whose first language is ASL does not have to translate English in her head while a train door closes. The trouble begins when the two versions do not say the same thing. A hearing passenger who notices a mismatch between the screen and the overhead voice can ask someone, or trust the voice. A Deaf passenger looking at a screen where the text says one thing and the signer says another has no third source. She has to guess which of two official messages, issued side by side by the same operator, is the true one.

That is what one Deaf rider who uses the train regularly recorded at CVG, and it is what Novara Consulting Group put in front of the Kenton County Airport Board in October 2026: AI-generated ASL on the train platform screens that contradicted the English text displayed beside it. The rider did not need a vendor evaluation, a word error rate or a model card to see the problem. She needed the signing to match the screen, and it did not. Every argument that follows, about baselines and comparators and procurement documents, is in the end an argument about whether that rider's experience counts as a failure or as an improvement on nothing.

A request to lower the bar

About thirty-five minutes into the fourth session of Innovate UK Business Connect's Accessible and Inclusive Journeys Innovation Cafe (https://iuk-business-connect.org.uk/events/accessible-and-inclusive-journeys-an-innovation-cafe/), held online on 27 October 2026 and recorded on Vimeo as AT&I Cafe 4 (https://vimeo.com/1241344794), Marcus, presenting for the sign language AI vendor Signapse, made a request that buyers of sign language AI will hear more often as the market matures. The panel had spent twenty minutes insisting that AI signing is not meant to replace the interpreting profession, and he agreed that this was true. The difficulty, as he saw it, was that interpreting had become the baseline against which the new technology was judged, and the vendors did not want to compete with it at all. The correction offered to transport operators was simple: the proper comparison is not AI signing

against an interpreter, it is AI signing against no access whatsoever. Making something accessible, he said, should itself be the baseline [6].

The argument deserves a careful hearing because it is persuasive, because it will be repeated in procurement rooms, and because it is wrong in a way that matters for anyone who signs a contract on the strength of it. It is not a claim about the quality of any product. It is a claim about which yardstick the buyer should pick up before measuring, and the yardstick proposed is one that no public body with an access duty is permitted to use.

What a baseline swap does

Every evaluation rests on a comparator, whether or not anyone names it. When a buyer asks whether a product is good enough, the unspoken second half of the question is good enough compared with what. Change the comparator and the same product, with the same error rate and the same comprehension results, can move from unacceptable to admirable without anything about it changing. This is why the choice of baseline is a governance decision rather than a technical detail, and why a vendor's proposal to change it should be treated with the same seriousness as a proposal to change the acceptance criteria themselves.

Measured against a qualified interpreter or a reviewed human-signed video, an AI signing system has to demonstrate accuracy, grammatical fidelity, intelligibility to the people it is meant to serve, and a tolerable failure mode when it gets something wrong. Measured against nothing, it has to demonstrate only that it exists. The second test is not a lower version of the first; it is a different test, one that a product cannot fail. What looks like a modest plea for fairness is in practice a request to be exempted from evaluation.

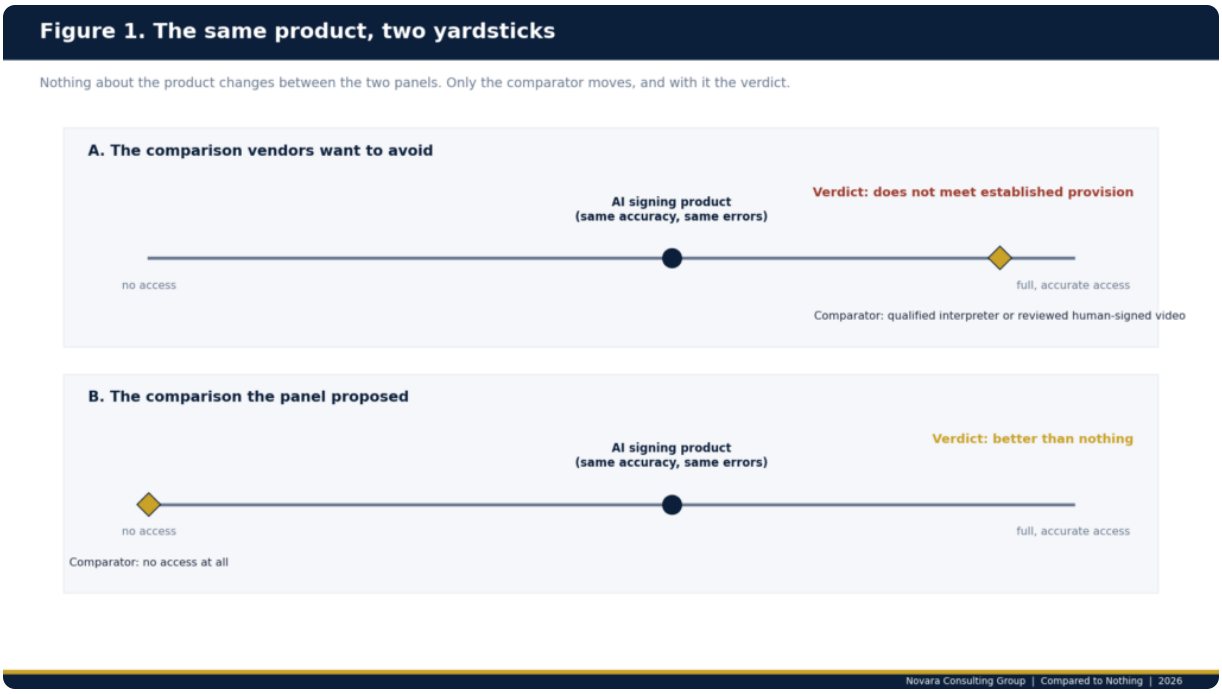


Figure 1. The same product, two yardsticks. Nothing about the product changes between the panels. Only the comparator moves, and with it the verdict.

The move works on operators for a reason that has little to do with sign language. It reframes a quality decision as a generosity decision. An operator who declines an AI signing product after comparing it with compliant provision has made a procurement judgment. An operator who declines it after being told the alternative is nothing has, in the vendor's framing, chosen to leave Deaf passengers with no access at all. Few procurement officers want to be the person who made that choice, and the framing is built to make sure they feel it. The rhetorical strength of the argument is precisely what should prompt a closer look at its premise.

The law has already chosen the baseline

The premise fails on the first contact with the statutes that govern the buyers being addressed. A transport operator in the United Kingdom or the United States is not choosing between AI signing and nothing, because nothing is not a state the law allows it to occupy. The baseline for a public body with an access duty is compliant provision, and the question a buyer must answer is whether the AI product meets or improves on that provision, not whether it improves on a void.

In the United States, a public entity such as an airport authority is bound by Title II of the Americans with Disabilities Act and its implementing regulation at 28 CFR 35.160, which requires the entity to furnish appropriate auxiliary aids and services where necessary to afford a person with a disability an equal opportunity to participate in and benefit from the service. The regulation requires the entity to give primary consideration to the aid or service requested by the individual, and it conditions the choice of aid on whether it is effective. Federally assisted transport facilities carry the parallel obligation under Section 504 of the Rehabilitation Act. An operator that currently provides nothing for Deaf passengers is not at a neutral starting point; it is in breach, and the breach is the operator's, not the vendor's problem to solve on the operator's behalf.

In the United Kingdom, the Equality Act 2010 imposes on service providers and those exercising public functions an anticipatory duty to make reasonable adjustments, including the provision of auxiliary aids, and the Public Sector Equality Duty requires public authorities to have due regard to the need to advance equality of opportunity for disabled people. Here too, the absence of any provision for signing passengers is not the baseline. It is the failure the duty exists to prevent.

Once this is seen, the real menu of options becomes visible. For a fixed, scheduled announcement at a gate, the compliant alternatives to AI signing include on-screen text, captioned audio, and pre-recorded human-signed video reviewed before publication. For a live, unscripted need, they include video remote interpreting and an on-site interpreter. These are the comparators. Some are cheaper than AI signing and some are more expensive, some are faster to deploy and some slower, and a buyer is entitled to weigh them on cost, speed, and reach. What a buyer is not entitled to do is pretend they do not exist in order to make the product under evaluation look like the only alternative to exclusion.



Figure 2. Where the law puts the baseline. A public body with an access duty is never choosing between AI signing and nothing; nothing is a breach, not a starting point.

The law has done this before

The baseline argument presents itself as a new problem created by a new technology, but American disability law has met this situation before and answered it the same way each time. When a cheaper, remote or automated alternative to established provision arrives, regulators have not asked whether it beats nothing. They have written down what it must do to count as provision at all.

The clearest precedent is video remote interpreting. When remote interpreting over video spread through hospitals, courts and government offices in the 2000s, it was sold on much the same terms now used for AI signing: on-site interpreters are expensive and hard to schedule, and a screen on a cart is available at once. The Department of Justice did not respond by treating VRI as an improvement over having no interpreter. In the 2010 revision of the Title II regulations it set performance standards that a public entity choosing VRI must meet [2]. The video must be real-time and full-motion over a dedicated high-speed connection, without lags, choppy, blurry or grainy images, or irregular pauses in communication. The image must be sharply delineated and large enough to show the interpreter's face, arms, hands and fingers, and the signer's, whatever their body position. The audio must be clear, and staff must be trained

to set up and run the system quickly. A VRI session that fails those standards is not a lesser form of compliance. It is not compliance.

The logic is worth holding onto because it transfers directly. The regulation did not compare VRI with nothing, and it did not even compare VRI with an on-site interpreter on every axis. It identified what effective communication requires in the medium being used and made those requirements a condition of using the medium. Applied to AI signing on a platform screen, the same reasoning asks what effective communication of a train departure or a safety instruction requires, and the first answer is obvious: the signed message must say what the written message says.

Captioning followed the same path. Broadcasters long argued that imperfect captions were better than none, and for years the rules required captions to exist without saying much about whether they were any good. In 2014 the Federal Communications Commission adopted caption quality standards addressing accuracy, synchronicity, completeness and placement [4]. Again the move was from presence to performance: a caption that is wrong, late or missing words was recognised as a failure of access rather than a partial success.

Air travel has its own version of the rule, and it is the closest to the facts at CVG. Under the Air Carrier Access Act regulations, a U.S. carrier must ensure that passengers who identify themselves as needing visual or hearing assistance have prompt access to the same information provided to other passengers at each gate, ticketing area and customer service desk, including information on flight safety, delays, schedule changes, boarding and gate assignments [3]. That obligation runs to airlines rather than airport operators, and the train platform at CVG is the airport's, so the rule does not decide the CVG complaint. But it shows plainly what standard federal transport regulation applies to information for Deaf passengers. The test is not some information, or better information than nothing. It is the same information. A signed message that contradicts the text beside it fails that test by definition, wherever it appears in an airport.

Wrong signing is not zero

The second flaw in the argument is arithmetical. The comparison to nothing assumes that the worst an AI signing system can do is fail to add value, so that any output at all is a net gain over silence. That assumption holds only if errors are harmless, and in a

transport setting they are not. A signed announcement that gives the wrong gate, the wrong platform, or the wrong safety instruction does not leave the Deaf passenger where they were. It moves them, with confidence, in the wrong direction. The passenger has no way to detect the error, because the whole point of the provision is that it stands in for information they cannot otherwise receive. Inaccurate access is therefore not a weaker form of access. It is misinformation delivered with the appearance of access, and it carries a cost that silence does not.

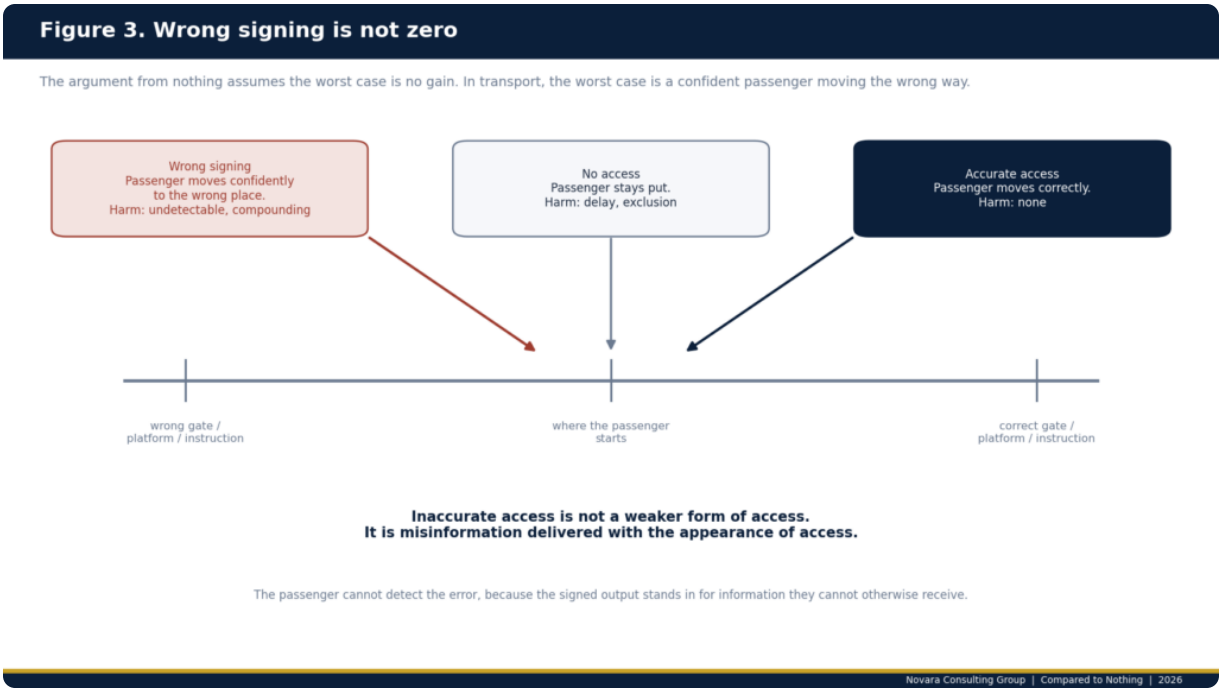


Figure 3. Wrong signing is not zero. The argument from nothing assumes the worst case is no gain; in transport, the worst case is a confident passenger moving the wrong way.

This is not a hypothetical. In October 2026, Novara Consulting Group filed a complaint under ADA Title II and Section 504 with the Kenton County Airport Board concerning AI-generated ASL announcements displayed on terminal train platform screens at Cincinnati/Northern Kentucky International Airport, where recorded signed content contradicted the English text on the same screens. One of the recordings came from a Deaf passenger who rides that train regularly. The airport subsequently suspended the signed train-screen announcements pending review. On the slide Signapse showed at the Innovate UK panel on 27 October, the day after that suspension, the same deployment appears in the vendor's airport table as providing ASL

on the boarding gate screen plus live tram departure and safety information, with the status given as live since 2023.

Figure 7. Signapse airport deployments as presented on 27 October 2026

Transcribed from the slide at 18:52 of the recording. Annotation in the right-hand column is NCG's, from its correspondence with the airport.

Airport	What Signapse provides	Status (as presented)	Known to NCG
Heathrow T5, with British Airways	Live BSL boarding announcements at five gates, via SignStream	Pilot complete, 2025	
CVG Airport, Cincinnati (US)	ASL on boarding gate screen, plus live tram departure and safety information	**Live since 2023**	**Train-screen ASL pulled pending review, 26 Oct 2026, after ADA complaint; matter with FAA**
Canadian airports, via Spoke	ASL versions of live announcements in Spoke's passenger app	Live, 2026	
GRR Airport, Grand Rapids (US)	Static ASL messaging, following a FLITE-funded live-board pilot	Static messages live	
MIA, Miami (US)	Real-time ASL departure information on airport screens	Pilot complete, 2025	

Source: vimeo.com/1241344794 at 18:52. Deployment claims are the vendor's; NCG has not independently verified rows other than CVG.

Novara Consulting Group | Compared to Nothing | 2026

Figure 7. Signapse airport deployments as presented on 27 October 2026, transcribed from the slide at 18:52. The right-hand annotation is NCG's, from its correspondence with the airport.

The slide also illustrates a second problem with the way deployments are described to buyers. The word live is doing two different jobs and the slide does not say which. It may mean that a live schedule feed triggers pre-built, pre-reviewed signed clips from a template library, which is how the vendor's streaming product has been described publicly. Or it may mean that spoken announcements are captured, recognised, translated, and rendered as ASL as they happen, which is a far larger technical claim and one that nothing in the vendor's public materials supports. A buyer needs to know which, because the two have different failure modes, different review requirements, and different answers to the question of who approved what the passenger saw. If the CVG content was templated, then signing that contradicts the screen text is not a translation error at all. It is a content, mapping, or change-control failure, and somewhere there is an approval record for the clip that was wrong.

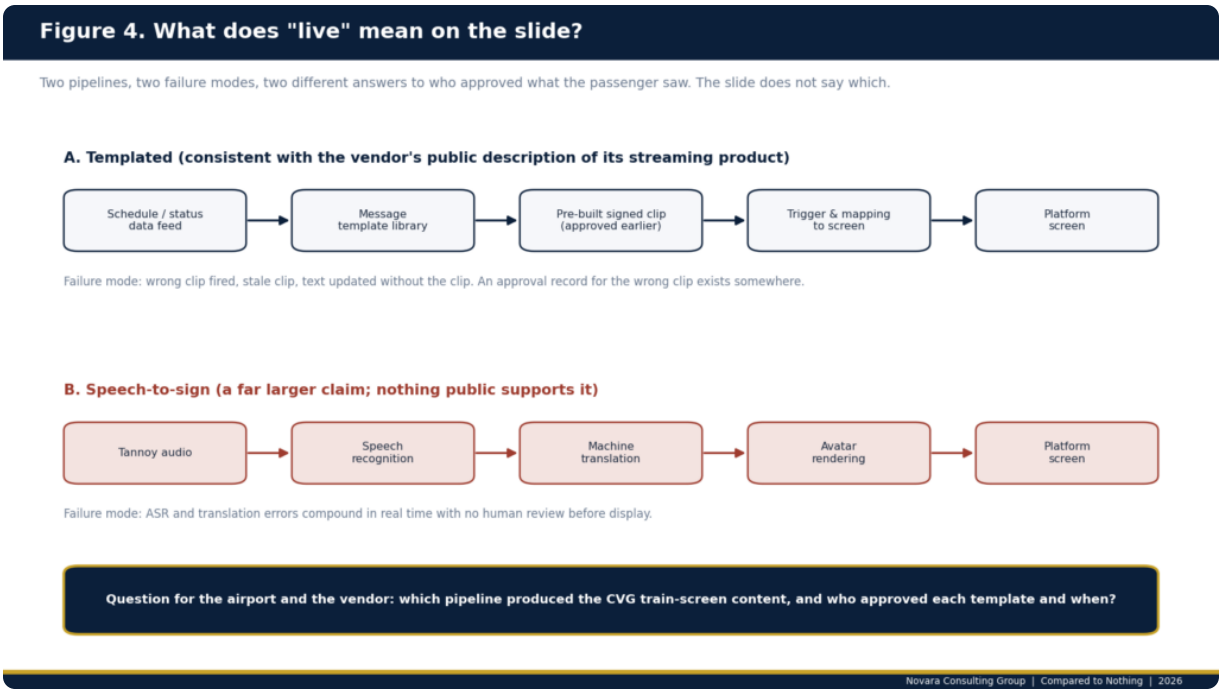


Figure 4. What does live mean on the slide? Two pipelines, two failure modes, two different answers to who approved what the passenger saw.

Either way, the argument from nothing does not survive the example. The operator did not have nothing before. It had English text on a screen, which is compliant provision for a fixed announcement, and it layered onto that text a signed version that contradicted it. The relevant comparison was never AI signing against a blank screen. It was AI signing against the accurate text already on the screen, and on that comparison the product subtracted value.

The community answered the comparator question in 2018

The comparator question did not need a transport panel in 2026 to raise it. Eight years earlier, the World Federation of the Deaf and the World Association of Sign Language Interpreters issued a joint statement on signing avatars that was addressed less to developers than to buyers [5]. Its stated concern was the way public authorities were deciding where and when to use avatars as a form of access to spoken or written content. The two federations cautioned against using avatars as a replacement for human signers, noted that the computerised products did not match the quality of trained and qualified interpreters and translators, and singled out the cases that worried them most: information that is live, complex or of significant importance to the lives of Deaf citizens.

A train departure and a safety instruction on an airport platform fall squarely within that last category. The information is time-bound, it governs where a person physically goes next, and getting it wrong has consequences that the passenger bears alone. The statement also made a point that the baseline argument skips over entirely: avatars provide one-way communication only, from spoken or written language into sign, and not back. A Deaf passenger who cannot make sense of a signed message on a screen cannot ask the screen what it meant. The channel that was supposed to provide access gives her no means of checking it.

Read against that statement, the Innovate UK panel's request looks less like a fresh argument and more like an attempt to move the decision out of the frame the Deaf community has already set. The international bodies representing Deaf people and sign language interpreters did not ask buyers to compare avatars with nothing. They asked buyers to recognise that the human signer is the standard for important information and to justify any departure from it. A vendor proposing a different baseline is not filling a gap in that guidance. It is asking procurement officers to set the guidance aside.

This matters for governance as well as for rhetoric. Deaf-led evaluation exists precisely because hearing buyers are poorly placed to judge whether signed output is accurate, intelligible and culturally appropriate, and are therefore the audience most easily persuaded that any signing is good signing. A comparator chosen by the vendor and accepted by a hearing procurement team is the arrangement most likely to let a contradiction like the one at CVG reach a platform screen, and stay there until a Deaf rider happens to record it.

What the vendor's own table shows

The deployments slide repays a closer reading than its presenter gave it, because it is the vendor's own account of its airport record, offered to an audience of potential buyers. Five rows appear. Two of them, Heathrow Terminal 5 with British Airways and Miami International, are described as pilots completed in 2025. One, Grand Rapids, is described as static ASL messaging following a pilot of a live board, which on the slide's own terms means the live element did not continue. One, the Canadian airports reached through Spoke's passenger app, is described as live in 2026 but sits inside a third party's app rather than on airport screens. That leaves a single row in which the vendor claims a live, on-airport, on-screen deployment that has run for more than a year: CVG, live since 2023.



That is the row that was suspended pending review on or about 26 October, the day before the slide was shown. Read with that fact in mind, the table describes a portfolio in which the longest-running live airport deployment on screens was the one a Deaf rider found contradicting its own text, and in which the remaining entries are completed pilots, a scaled-back static service and an in-app feature. None of this proves anything about the quality of the other deployments, and NCG has not examined them. It does show that an audience shown the slide would have come away with an impression of established, ongoing live service that the vendor's own row-by-row descriptions, read carefully, do not support, and that the one row that did support it had just been taken off the screens.

Buyers should read every vendor deployment table this way. The words pilot, static and live carry very different evidentiary weight, and a slide that sets them side by side under a single heading invites the eye to treat them as equivalent. A procurement team that asks, row by row, what is running now, on whose screens, with what content approval process and with what complaint history, will learn more from a deployments slide than its author intended.

The record, in order

The sequence of events is short, each step is independently verifiable, and the two sources that matter most, the airport's own correspondence and the vendor's own recorded presentation, were produced without reference to each other. That is what makes the discrepancy between them useful to a regulator. Nothing here depends on NCG's characterisation; it depends on dates and on what each party said in its own words.

The record, October 2026

Date	Event	How a regulator can verify it
12 Oct 2026	NCG files an ADA Title II and Section 504 complaint with the Kenton County Airport Board's Title VI/ADA coordinator over AI-generated ASL on CVG terminal train platform screens that contradicts the screen text; two recordings attached; the Board acknowledges receipt the same day	Board's complaint file; NCG's filed complaint and acknowledgement; the two screen recordings
12 Oct 2026	NCG submits a Kentucky Open Records request to the Board for the Signapse contract and procurement file and testing records and complaint history	Board's open records log
19 Oct 2026	The Board declines the records request on residency grounds under KRS 61.870(10) and invites resubmission under another category; the contract and testing records remain unproduced	Board's written denial
On or about 26 Oct 2026	The Board informs NCG that the signed train-screen announcements have been pulled pending review and forwards the matter to the FAA; NCG files its own complaint with the FAA Office of Civil Rights the same day	Board's correspondence to NCG; FAA Office of Civil Rights intake
27 Oct 2026	Signapse presents at session 4 of the Innovate UK Business Connect Accessible and Inclusive Journeys Innovation Cafe (13:00 to 13:45 UK time); at 18:52 the deployments slide lists CVG as live since 2023 with live tram departure and safety information; at 34:46 to 35:14 Marcus of Signapse argues that AI signing should be judged against no access	Vimeo recording at the timestamps stated; session date on the Innovate UK Business Connect event page
30 Oct 2026	The recording is uploaded publicly to Vimeo by Innovate UK Business Connect	Vimeo upload metadata

Source: NCG correspondence with the Kenton County Airport Board; Innovate UK Business Connect event page and Vimeo recording

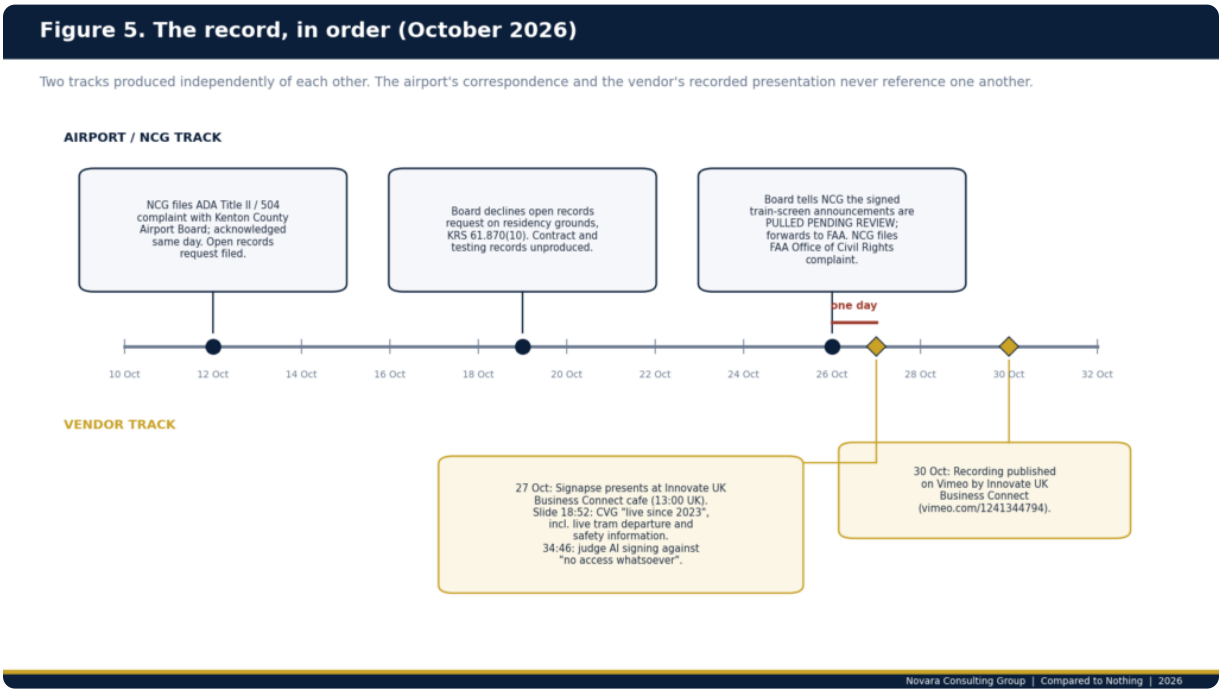


Figure 5. The record, in order. Two tracks produced independently of each other; the airport's correspondence and the vendor's recorded presentation never reference one another.

Four points follow from the record, and each is a question the FAA can put to the parties rather than a conclusion it has to take on trust.

First, the vendor presented the CVG train-screen deployment to a public audience as live, and as having been live since 2023, at a time when the airport had already told the complainant that the signed train-screen announcements were pulled pending review. Either the vendor had not been told of the suspension by its customer, or it had been told and presented the deployment as live anyway. Both answers bear on the adequacy of the arrangements between the airport and its vendor, and the FAA can ask which is the case.



Figure 6. Two statements, one day apart. Each party in its own words; the question for the FAA is which explanation for the gap applies.

Second, the slide describes the CVG content as including safety information. If safety messaging passed through the same pipeline that produced signed announcements contradicting the screen text, then the complaint concerns safety-critical content, not only wayfinding, and the review the airport has undertaken should be scoped accordingly.

Third, the word live on the slide is undefined. The FAA can ask the airport and the vendor whether the signed output was generated from live spoken announcements or rendered from pre-approved templates triggered by a data feed, and if the latter, who approved each template and when the templates were last reconciled against the screen text. The answer determines whether the contradictions were a translation failure or an approval and change-control failure, and therefore who at the airport signed off on content that was wrong.

Fourth, the baseline argument made at 34:46 is the argument against which the CVG facts should be read. Signapse′s presenter asked that AI signing be compared with nothing. At CVG it was compared, by a Deaf passenger with a camera, with the accurate English text on the same screen, and it lost. The vendor′s proposed yardstick

would have made that result invisible. The airport's actual duty under 28 CFR 35.160 made it reportable.

Who holds the duty when the vendor is wrong

One consequence of the baseline argument is easy to miss because it is buried in its structure. If AI signing is framed as a voluntary improvement over nothing, then a failure of the AI looks like the failure of a gift, regrettable but nobody's breach. The law does not see it that way. A public entity may not discriminate on the basis of disability directly or through contractual, licensing or other arrangements [1]. The operator's duty to communicate effectively does not pass to the vendor when the operator buys a product, and it does not shrink to whatever the product happens to deliver. When the signed message on the screen is wrong, the operator has communicated something wrong to a Deaf passenger, through a contractor it chose.

That allocation of responsibility changes what a careful operator needs from its vendor, and it explains why the gap between 26 and 27 October is more than an embarrassment. An operator that carries the duty needs to know, continuously, what its contractor is putting in front of passengers, how the contractor checks it, and what happens when something goes wrong. It also needs the contractor to describe the deployment to others accurately, because the operator's name is on the deployment. A vendor that presents a suspended deployment as live to an international audience of transport operators is, at minimum, out of step with its customer. Whether that is because nobody told the vendor, or because the vendor did not consider the suspension worth mentioning, the operator has a contract management problem either way.

The records that would settle these questions are not exotic. They are the contract and its statement of work, any acceptance testing or quality assurance plan, the content approval process for signed messages, the change-control procedure for keeping signed content aligned with screen text, the incident and complaint log, and any correspondence between the airport and the vendor after the complaint was filed. These are the same records NCG asked for under the Kentucky Open Records Act on 12 October and that remain unproduced after the residency denial of 19 October. They are also records that a federal civil rights reviewer can request directly, without any residency requirement, and they are the fastest route to answering the four questions set out above.

The quiet part

There is a candour in the panel's framing that buyers should notice. Vendors in this market routinely say that their products are not intended to replace interpreters, and the statement is usually offered as reassurance to the Deaf community and the interpreting profession. On this panel it was paired with a complaint: that the technology keeps being compared with interpreting, which the vendors do not wish to compete with. Read together, the reassurance and the complaint are the same sentence said twice. A product that cannot be compared with the established provision is a product that does not perform as well as the established provision, and the request not to be compared is an admission of the result.

None of this is an argument that AI signing has no legitimate use. There are settings in which the established provision is in fact absent, unreachable in the time available, or disproportionate to the message, and in those settings a well-governed system with a known accuracy profile and a reviewed content library may be the right tool. But the way to establish that is to name the comparator honestly and show that the product meets it in the defined setting, not to replace the comparator with a void and declare victory over emptiness. A vendor confident of its product's performance would welcome the harder comparison. The request to be spared it tells the buyer most of what the evaluation would have found.

What a fair comparison would look like

Rejecting the comparison to nothing is only half the work. The other half is saying what an honest evaluation of AI signing in a transport setting would involve, so that operators who want to use the technology responsibly have something to hold it to. None of the elements below is novel, and none requires a buyer to become an expert in machine learning. They require the buyer to treat signed output as communication that a person will act on, which is what it is.

The evaluation starts by sorting messages by consequence. A welcome message, a scheduled departure, a gate change, a delay and an evacuation instruction do not carry the same risk, and a single pass or fail for the whole system hides the difference. For each class, the operator states the compliant provision it would otherwise use, which for most fixed transport messages is accurate on-screen text, often supplemented by a reviewed, pre-recorded human-signed clip. The AI output is then tested



against that provision, message class by message class, with the highest-consequence classes held to the strictest standard or excluded altogether.

The testing itself has to involve the people the output is for. Fluent Deaf signers, and for American deployments that means fluent users of ASL, should view the signed output and report what it says, without seeing the text it was generated from. Their account is then compared with the text. This is the simplest possible test of the property that failed at CVG, and it measures meaning rather than appearance. Errors should be classified rather than merely counted, because a signed message that is stiff but correct is a different kind of problem from one that is fluent but says the wrong platform. Content errors, where the signed message conveys different information from the text, should be treated as disqualifying for any message class that governs where a passenger goes or what a passenger does.

The evaluation does not end at acceptance. Templated systems drift whenever the underlying text changes and the signed library does not, and live systems drift whenever the language they are given moves outside what they were tested on. An operator therefore needs a change-control rule that no screen text is altered without the corresponding signed content being reviewed, a routine sampling check by Deaf reviewers on what is actually being displayed, a visible and accessible way for Deaf passengers to report a mismatch, and a commitment to pull signed content immediately when one is reported, with the text left in place. None of that is expensive compared with the cost of the deployment itself, and all of it would have caught the CVG contradiction before a passenger had to.

Finally, the evaluation is written down before the vendor presents, and the vendor's evidence is assessed against it rather than the other way round. That single procedural step is what defeats the baseline swap. A vendor cannot move a comparator that the buyer has already fixed in its own documents.

Writing the comparator into the evaluation

The defence against a baseline swap is procedural and inexpensive. A buyer who names the comparator in the evaluation documents before any vendor presents has removed the vendor's ability to choose it. The comparator should be stated per message type rather than per contract, because the compliant alternative to a scheduled gate announcement is not the same as the compliant alternative to an emergency instruction, and a product that is adequate for the first may be unsafe for the second.

For each message type the evaluation should record the provision the operator currently makes or is obliged to make, the accuracy and comprehension evidence the vendor offers against that provision, and the consequence to the passenger of a wrong output.

Three questions belong in every request for proposal in this market and in every contract review of an existing deployment.

1. Against which existing provision is the product being evaluated for each message type, and who chose that comparator?
2. Is the signed output generated from live speech, or rendered from pre-scripted text or data-feed templates? If templated, who approved each template, and what is the change-control process when the corresponding text changes?
3. What is the documented failure mode when the output is wrong, how would a Deaf passenger detect it, and who bears the duty when they cannot?



Figure 8. The baseline-selection test for buyers. Name the comparator before the vendor can.

A vendor who can answer these plainly has nothing to fear from the interpreter comparison. A vendor who responds by proposing a different baseline has answered the first question in a way the buyer should write down.

Something is not better than nothing

The sentence that opened this piece will be heard in procurement rooms in both countries for years, because it is short, humane, and flattering to the buyer. It should be answered with an equally short reply. The operator was never choosing between something and nothing. It was choosing between the provision the law requires and a product that asks not to be measured against it. A buyer who accepts the proposed baseline has not lowered the bar for the vendor; it has lowered the bar for itself, and the Deaf passenger who follows a confidently wrong announcement to the wrong platform is the one who pays the difference.

SOURCING NOTE

The panel discussed here is session 4 of the Innovate UK Business Connect series Accessible and Inclusive Journeys: An Innovation Cafe, part of its Accessible and Inclusive Transport programme, held online on 27 October 2026 from 13:00 to 13:45 UK time according to the event page (https://iuk-business-connect.org.uk/events/accessible-and-inclusive-journeys-an-innovation-cafe/), and published by Innovate UK Business Connect on Vimeo on 30 October 2026 under the title AT&I Cafe 4 2710 at https://vimeo.com/1241344794. The baseline argument is made by Marcus of Signapse between 34:46 and 35:14 of the recording; the airport deployments slide appears at 18:52 and is credited on screen to Signapse. The CVG complaint, the open records denial, and the airport's suspension of the signed train-screen announcements are matters of NCG's own correspondence with the Kenton County Airport Board in October 2026, and copies are available to the FAA on request. Statutory references are to 28 CFR 35.160 under ADA Title II, Section 504 of the Rehabilitation Act of 1973, KRS 61.870(10), and sections 20, 29, and 149 of the Equality Act 2010. Characterisations of how the vendor's streaming product generates output are drawn from the vendor's public descriptions and should be read as NCG's understanding rather than as confirmed technical fact.

References

1. U.S. Department of Justice. 28 CFR 35.130, General prohibitions against discrimination. Code of Federal Regulations, Title 28, Part 35 (ADA Title II).

2. U.S. Department of Justice. 28 CFR 35.160, Communications: General, including paragraph (d), Video remote interpreting (VRI) services. Code of Federal Regulations, Title 28, Part 35 (ADA Title II), as revised 2010.
3. U.S. Department of Transportation. 14 CFR 382.53, What information must carriers give individuals with a vision or hearing impairment at airports? Code of Federal Regulations, Title 14, Part 382 (Air Carrier Access Act).
4. Federal Communications Commission. Closed Captioning of Video Programming; Telecommunications for the Deaf and Hard of Hearing, Inc. Petition for Rulemaking. Report and Order, Declaratory Ruling, and Further Notice of Proposed Rulemaking, FCC 14-12, February 2014.
5. World Federation of the Deaf and World Association of Sign Language Interpreters. WFD and WASLI Statement on Use of Signing Avatars. Helsinki and Melbourne, 14 March 2018, updated 14 April 2018. https://wfdeaf.org/wp-content/uploads/WFD-and-WASLI-Statement-on-Avatar-FINAL-14032018-Updated-14042018.pdf
6. Innovate UK Business Connect. Accessible and Inclusive Journeys: An Innovation Cafe, Session 4, 27 October 2026. Event page: https://iuk-business-connect.org.uk/events/accessible-and-inclusive-journeys-an-innovation-cafe/ Recording, AT&I Cafe 4 2710, published 30 October 2026: https://vimeo.com/1241344794

How to cite

Grizzle, H. M. (2026, November 1). Compared to Nothing: The Baseline Swap in Sign Language AI Sales. Novara Consulting Group. https://www.novaracg.com/2026/11/01/compared-to-nothing-the-baseline-swap-in-sign-language-ai-sales/