This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/10/25/never-signed-or-signed-wrong-what-two-airports-and-a-rail-foi-teach-buyers-of-ai-signing/.

Never Signed, or Signed Wrong: What Two Airports and a Rail FOI Teach Buyers of AI Signing

Heather Grizzle · October 25, 2026

A CVG signer contradicted its screen, UK rail never signs spoken updates, and Grand Rapids scoped it right. What buyers of AI signing must require.

Contenuti

1. Two screens, two failures
2. One design decision behind both failures
3. A buyer who scoped it
4. Why disruption is where it breaks
5. A station count is not an access claim
6. The third failure: never seen
7. What the vendor and the airport have said so far
8. What buyers should require
9. The same information, at the same time
10. Riferimenti

IN BREVE

- At Cincinnati/Northern Kentucky International Airport, two recordings published this month show an AI signer telling Deaf travelers something different from the text on the same screen.
- At Gerald R. Ford International Airport in Grand Rapids, the same supplier's signing runs only where the messages repeat and do not depend on real-time translation. That is a buyer matching the deployment to what the system can actually carry.
- Freedom of Information replies obtained by Tim Scannell show the opposite gap on UK rail: announcements spoken over the public address system are not automatically signed.
- An automated signed announcement can fail a Deaf passenger three ways. It can be never signed, signed wrong, or hung where the passenger never sees it. All three are decided before installation, by the buyer.
- The fixes belong in the contract: a message-by-message coverage statement, a named input, a reconciliation check that fails closed, placement reviewed by Deaf and DeafBlind users, and a published account of what is not covered.

Two screens, two failures

The first thing I check at any airport screen is where it hangs. I am DeafBlind and I cannot see above eye level, so a signer mounted near the ceiling is a signer I will never see. Only after that do I get to the question this piece is about: when the screen is where I can see it, does it tell me what everyone else is being told?

At Cincinnati/Northern Kentucky International Airport this month, it did not. On 8 October 2026, a video posted to LinkedIn showed Max, the AI-generated ASL signer Signapse supplies, on a live screen on the airport's terminal train. The text on the screen read "Next train arriving in 5 minutes for A & B Gates." The signer said the train would arrive in about two to three minutes and that it served baggage claim and ticketing, which describes the train going the other way [2]. Four days later, a Deaf traveler who rides that train regularly supplied a second recording from the same screens, in which the signer named a different concourse from the one the train serves

[3]. In both recordings, the hearing traveler standing beside a Deaf traveler got the right information, and the Deaf traveler did not.

Across the Atlantic, Tim Scannell asked British rail operators under the Freedom of Information Act what the BSL on their station screens actually covers [1]. Network Rail, Govia Thameslink Railway and West Midlands Trains each told him some version of the same thing: an announcement spoken over the public address system does not automatically become BSL. West Midlands Trains, which uses Signapse at 24 stations, said its recorded avatar currently covers non-stopping trains passing through. Scannell is careful to call the risk to Deaf passengers a potential consequence rather than evidence that anyone was harmed, and that care is right. Nothing on those screens is reported to contradict the text. What the replies describe is silence.

Read together, the two bodies of evidence show the two ways an automated signed announcement can fail the person it exists for. It can never be signed, or it can be signed wrong.

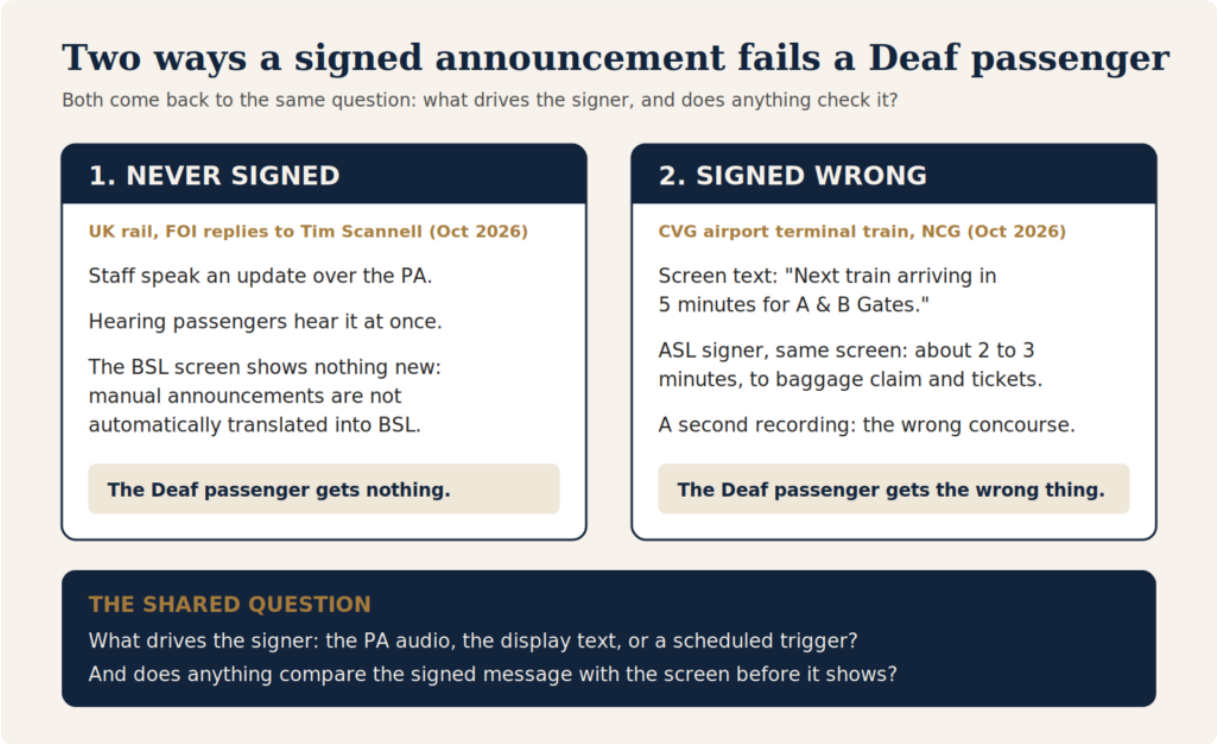

Two failures, one shared question: what drives the signer, and does anything compare it with the screen before it shows?

One design decision behind both failures

Signapse describes its transport product as a library of pre-recorded signed clips paired with announcement types, which software assembles into a signed message and updates from live transport data such as destinations, times and platforms [2]. West Midlands Trains describes the same thing from the buyer's side: a recorded avatar. A system built this way can sign only what reaches it as data. A sentence a member of staff speaks into a microphone never becomes data, so it never becomes sign. That is the gap the rail replies describe.

The same architecture explains how the second failure can happen. If the screen text and the signed clip are produced by separate paths, or the mapping from one to the other selects the wrong clip, the two can disagree, and the screen will show both at once. Which of these happened at CVG is not public. The signed clip there may follow the public address audio, the display text, or a scheduled trigger, and only Signapse and the airport can say which [3].

So the two failures are not two unrelated defects. They are two outcomes of one design decision: what the signer listens to. Omission happens when the signer listens to a feed that misses things. Contradiction happens when the signer and the screen listen to different things, or when nothing checks that they agree before both reach a passenger. A buyer who cannot answer the first question cannot know which failure they have bought.

A buyer who scoped it

The same supplier's signing has been running for some time at Gerald R. Ford International Airport in Grand Rapids, Michigan, and the way it got there is worth more to a buyer than any station count. When I asked how the airport's FLITE innovation program had judged its Signapse pilot, the program explained that Signapse connected its software to the airport's flight data feed and showed signed schedule information on a screen beside the existing departure boards. The program put a QR-code survey next to the screen and brought a representative from Deaf and Hard of Hearing Services on site to give feedback [4].

Then the airport made the decision that matters. Signapse was at that point still working toward real-time translation, so the airport placed the signing at the TSA PreCheck checkpoint and at the gates, where the messages repeat and do not depend on

translating something new in the moment. The signing is still running in both places [4].

That is a buyer matching the deployment to the kind of message the system can reliably carry, and keeping it away from the kind it cannot. It is the decision the rest of this piece argues for, made in practice by an airport that was not asked to make it.

The program also told me one of its biggest lessons was how much it meant for a Deaf passenger to walk in and see their own language. I believe that completely. The first time I saw ASL on a public screen, I felt it too. But recognizing your language and understanding the message are two different things, and only the second one gets you to the right gate. A buyer who measures the first and assumes the second has measured the wrong thing.

Why disruption is where it breaks

What follows is NCG's reading of the replies rather than a measured finding. Scheduled information is the easiest content to sign, because it is known in advance and already sits in the timetable data. Manual announcements cluster where reality runs ahead of the system: a late train, a platform swap, a fault, an incident, an evacuation. If signed coverage depends on what has already entered the operator's data systems, then it is strongest for the messages a Deaf passenger could already read on a departure board, and weakest for the ones nobody can read anywhere until a person says them aloud.

That is the opposite of where accessible information matters most. The routine announcement a Deaf passenger misses costs a few minutes. The unplanned one can decide whether they board the right train, wait on the right platform, or leave a station when told to. The British Deaf Association finding Scannell reports, gaps in emergency and unexpected journey information, points at the same place.

A station count is not an access claim

Operators reported their BSL by the number of stations, because that is the unit a rollout is planned and announced in. It is the wrong unit for an access claim. A station count describes where hardware is installed. Equal access is a property of messages: which kinds of announcement are signed, how soon after the event, whether the si-

gned version says what the screen and the speaker say, and whether the person it is for can see it at all.

What a station count tells you	What a Deaf passenger needs to know
Where the screens are	Which announcements those screens sign
How many stations have them	How long after an event the signed message appears
That a supplier was contracted	Whether the signed message matches the text beside it
That screens were installed	Whether they hang where a Deaf or DeafBlind passenger can see them
That the service exists	What is not covered, and where to go for it

None of the right-hand column is visible from a press release, and most of it is not visible from the platform either. A Deaf passenger looking at a screen cannot tell whether an absent update means nothing has changed or means the change was spoken and never signed. Nor can they tell a correct signed message from a fluent, confident wrong one. The better the signing looks, the more completely the error is trusted [2].

The third failure: never seen

Both failures above assume the passenger can see the screen. Many cannot. Wayfinding screens are usually hung high so a crowd can read them over each other's heads. That works for text a hearing traveler only needs to glance at. It works badly for a signer, a small moving figure whose hands, face and body all carry meaning, and it does not work at all for someone like me with no vision above eye level. Low vision, field loss, a wheelchair or simply a crowd standing in front of the screen produce the same result: the message was signed, correctly, to no one.

Deployment announcements report how many screens went up. They rarely say how high, how large the signer appears on them, or whether a single Deaf or DeafBlind traveler stood in front of one before it went live. Those are the questions that decide whether the signing reaches anybody, and like the other two failures, they are settled at installation, by the buyer.

What the vendor and the airport have said so far

On 12 October, NCG put one question to Signapse before publishing on the second CVG recording: what input drives the signed announcement on those screens. On 13 October, Signapse's chief executive replied, in full: "CVG are investigating the issue and we will be in touch with an update once it is resolved." That is a prompt and courteous reply, and NCG will publish any update in full. It does not yet answer the input question, and the input question is the one that decides where the fault sits and what control would fix it.

Signapse's own published governance material sets a higher bar than either recording meets, and its chief executive has argued publicly that higher impact and less preparation time call for more human oversight and less AI [2]. Live wayfinding and disruption announcements are exactly that combination.

The same day, I filed a complaint under Title II of the Americans with Disabilities Act and Section 504 with the airport's ADA coordinator, asking that the signer come off live wayfinding screens until something checks it against the text beside it. The airport acknowledged the complaint the same day. I also asked the Kenton County Airport Board, which runs CVG, for the Signapse contract, the procurement file, testing records and complaint logs under the Kentucky Open Records Act. On 19 October the Board declined the request because NCG is not a Kentucky resident, which the Act requires, and invited a new request under one of the categories it recognizes [5].

That is the law working as written. It also means the documents that would answer the input question sit in a public body's files, and the people most likely to read them critically, Deaf travelers and the specialists who evaluate these systems, are shut out unless they happen to live in the right state. A buyer does not have to wait for a records request to settle this. It can publish the coverage statement itself.

What buyers should require

All three gaps can be closed, and the place to close them is the solicitation and the contract, before a screen goes up. Contracts written before AI signing existed usually say nothing about it. In September, Minnesota's Department of Administration told NCG that the state's enterprise signing contracts predate AI and that it intends to review those contracts and write the next signing solicitation with state subject matter

experts [6]. The terms below are the kind of language that review needs, and they apply to any operator buying automated signing for passenger information.

MUST The provider MUST state, for each deployment, which message classes are signed and which are not: scheduled departures, passing trains, delays, cancellations, platform changes, route changes, manual announcements, and safety and evacuation instructions.

MUST The provider MUST name the input that triggers each signed message, and state whether it is the same input that writes the text on the screen.

MUST The system MUST compare each signed message with the displayed text before showing it, and MUST suppress the signed message whenever the two disagree on a time, destination, platform or instruction. Showing a Deaf passenger nothing is safer than showing them the wrong train.

MUST NOT The deployment MUST NOT be presented to passengers or the public as equivalent to spoken announcements while any message class is unsigned. The screen itself MUST say, in sign and in text, what it does not cover and where to get it.

MUST The provider MUST report the time between an event and its signed message, for each message class.

MUST Safety and evacuation instructions MUST reach Deaf passengers by a stated accessible route, whether or not the signed screens carry them.

MUST Screen placement, mounting height and the size of the signer on screen MUST be reviewed on site with Deaf and DeafBlind users before go-live, and the result MUST be recorded.

SHOULD The operator SHOULD keep a human route, through staff and remote interpreting, for every message class the screens do not cover, and SHOULD say on the screen how to reach it.

MUST The deployment MUST have an incident channel a Deaf passenger can actually use, and every mismatch reported through it MUST be logged and reported to the buyer.

SHOULD The buyer SHOULD publish the coverage statement, the named input and the verification method, so that the people the system serves can check the claim without a records request.

None of these terms asks a supplier to stop selling signed announcements, and none assumes automated signing is unacceptable. Grand Rapids shows the technology can be placed honestly. The terms ask that the claim made for it match what it does. That is the reason the SLAT Index evaluates sign language AI at the level of the deployment rather than the product: the same system can be adequate on one screen and misleading on another [7].

The same information, at the same time

Scannell closes his post with the right standard: Deaf passengers need the same information, at the same time. The rail replies show what happens to "at the same time" when information is spoken and never signed. The CVG recordings show that "the same information" cannot be assumed even when something is on the screen. Grand Rapids shows that a buyer who knows what the system can carry can keep it to that. I would add one clause, because I learned it before I learned anything about AI: the

same information, at the same time, on a screen I can see. The only point at which a buyer can guarantee all of it is the contract signed before the screen goes up.

Riferimenti

1. Scannell, T. (2026, October). BSL at railway stations: The same information, at the same time? [LinkedIn post]. https://lnkd.in/p/gqx-FCwH
2. Grizzle, H. M. (2026, October 8). The sign said something else. Novara Consulting Group. https://www.novaracg.com/2026/10/08/the-sign-said-something-else/
3. Grizzle, H. M. (2026, October 12). The check is still not there. Novara Consulting Group. https://www.novaracg.com/2026/10/12/the-check-is-still-not-there/
4. FLITE program, Gerald R. Ford International Airport. (2026, October 13). Correspondence to Novara Consulting Group regarding the Signapse ASL project.
5. Kenton County Airport Board. (2026, October 19). Response to Kentucky Open Records Act request from Novara Consulting Group regarding Signapse records.
6. Doran, A. (2026, September 15). Correspondence from the Minnesota Department of Administration, Office of State Procurement, to Novara Consulting Group regarding AI-based signing on state language services contracts.
7. Grizzle, H. M. (2026). SLAT Index Evaluation Standard (Version 2.0). Novara Press. https://doi.org/10.5281/zenodo.21537271

Come citare

Grizzle, H. M. (2026, October 25). Never Signed, or Signed Wrong: What Two Airports and a Rail FOI Teach Buyers of AI Signing. Novara Consulting Group. https://www.novaracg.com/it/2026/10/25/never-signed-or-signed-wrong-what-two-airports-and-a-rail-foi-teach-buyers-of-ai-signing/