

THE FIFTH PARAMETER | Issue 010 | Conformance Is Not Comprehension

Heather Grizzle · July 12, 2026

Contents

1. What the assessments show
2. The absence is structural
3. What a serious requirement looks like
4. The position

A procurement officer evaluating a signing avatar system has instruments for nearly everything except the thing that matters. There are established questions for encryption at rest and in transit, for data residency, for subprocessor disclosure, for uptime guarantees, for indemnification, for conflict minerals. There is a Voluntary Product Accessibility Template, and the vendor will supply one. There is a Section 508 checklist, and the vendor will satisfy it. At the end of this process the file is complete, the scoring rubric is filled, the award is defensible on audit, and no one at any point has asked whether a Deaf person watching the system's output understands what it is saying.

This is not an oversight by individual buyers, who are generally conscientious and working within the tools they have been given. It is a property of the tools. Accessibility conformance documentation certifies the presence of an alternative channel. It does not certify that the channel carries meaning. A VPAT can accurately record that signed content is provided alongside spoken content, and that record will be true, and it will tell the institution nothing whatever about whether the signing is intelligible to the population it exists to serve. The form has a line for whether the alternative was offered. It has no line

for whether it worked.

The consequence is a quiet inversion. Where an instrument exists, absence of evidence registers as a deficiency and gets flagged. Where no instrument exists, absence of evidence registers as nothing at all. The linguistic question is not answered badly; it is not posed, and its non-posing leaves no trace in the file. Silence is then read downstream as satisfaction, because a completed evaluation with no adverse findings looks identical to a completed evaluation in which the adverse finding could not be recorded.

What the assessments show

Novara Consulting Group has now conducted structured assessments of seven sign language AI vendors under the SLAT Framework: Kara Technologies, SignAvatar, ChatSign, Silence Speaks, Sign-Speak, Signvrse, and GoSign.AI. These are not equivalent products. They differ in architecture, in maturity, in the seriousness of their engagement with Deaf communities, and in their candor. Some are demonstrably better governed than others. Several employ Deaf staff, and a few have Deaf ownership or Certified Deaf Interpreters positioned in the production loop, which is a real distinction and is scored as one. Not one of them reached High Trust.

The binding constraint was the same in every case, and it was not any of the things a procurement file would have caught. In no instance was there independent validation of signed output with Deaf users: no comprehension testing conducted by a party without a commercial interest in the result, reported with its failure conditions intact, across the range of dialect, age, register, and language fluency that any public deployment will actually encounter. Vendors publish accuracy figures. Accuracy figures are generated by the vendor, against the vendor's reference data, on the vendor's terms, and they measure the distance between an output and a target rather than the distance between an output and understanding. The two are not the same quantity and no amount of the first

substitutes for the second.

A related pattern recurs across the snapshots.

The marketing surface consistently outruns the stated technical maturity, and the divergence is visible in the vendors' own documentation if one reads it against itself. Deployment-speed claims

appear without the conditions under which they were achieved. Systems that in fact operate along two distinct paths, one heavily human-mediated and one substantially automated, are described publicly in language that collapses the distinction. This is unremarkable behavior for a company raising capital. It is disqualifying as an evidentiary basis for a public contract, and the reason it survives contact with procurement is that nothing in the process is designed to test it.

The form has a line for whether the alternative was offered. It has no line for whether it worked.

The absence is structural

It would be a mistake, and an unproductive one, to read this as a vendor indictment.

Vendors produce the evidence that buyers demand. Procurement asks for SOC 2, so SOC 2 exists. Procurement asks for penetration test summaries, so penetration test summaries exist. Procurement has never once, in any instrument this publication has been able to locate, asked a sign language AI vendor to demonstrate that Deaf users comprehend its signed output under conditions resembling deployment. So that evidence does not exist. The market has supplied precisely what it was asked for.

This locates the lever on the buyer's side of the transaction, which is the argument this publication has been making since its founding and which is worth stating without hedging.

An RFP is not an administrative document. It is a specification for an evidence base.

Whatever the RFP requires, the sector will build; whatever it omits, the sector will not build,

not out of bad faith but because unfunded evidence does not get produced. Every requirement an institution declines to write is a study that will never be conducted. Procurement is the mechanism by which an institution's tolerance for uncertainty becomes the field's research agenda.

What a serious requirement looks like

The requirement is not complicated to state, though it is expensive to satisfy, and the expense is the point. Comprehension of signed output should be evaluated with Deaf participants who were not involved in building the system, spanning regional variation, age, and the range of language fluency the deployment will meet in practice, including the users who most need the system to work and are least equipped to compensate for its failures. The evaluation should be conducted by a party with no financial stake in its outcome. It should report where the system fails, under what conditions, and with what consequences for the user, because a system's failure modes are more informative for a buyer than its aggregate performance and are the first thing omitted from a vendor deck.

NOT MET AND NOT EVALUATED

These are different findings. Collapsing them punishes vendors who are honest about their maturity and rewards those who are not.

And institutions need a category, on the form, in the scoring, for evidence that does not yet exist. Not met and not evaluated are different findings. Collapsing them serves everyone poorly: it punishes vendors who are honest about their maturity and rewards those who are not, and it leaves the buying institution unable to say, in writing, the most accurate sentence available to it, which is that no one has checked.

The position

None of this is an argument against sign language AI, and readers looking for one will not find it here. The technology is early, some of it is serious, and some of it will likely become appropriate for public deployment. The claim is narrower and harder to evade: at present there is no procurement instrument in general use capable of distinguishing a system Deaf people understand from one they do not, the vendors have not been asked to supply that evidence and accordingly have not, and institutions are signing contracts in which the central question of the transaction is nowhere represented in the file.

A technology that has never sat the examination cannot be said to have passed it. Until the form has a line for that finding, it will keep going unrecorded, and the record will keep showing that everything was in order.

The Fifth Parameter is an independent critical publication on sign language AI, published by Novara Consulting Group.