

NCG Insight | When Accessibility Is Built Last, Risk Is Built In

Heather Grizzle · April 16, 2026

There is a familiar pattern in technology. A tool is introduced, framed as progress, and positioned as solving a long-standing barrier. The narrative moves quickly, often before the structure behind it is fully examined. By the time questions start being asked about impact, the foundation has already been built, decisions have already been made, and the people most affected are left responding to a system they did not help design.

Sign language (SL) AI is following that pattern. Over the past year, development has accelerated across recognition systems, translation tools, and avatar-based communication models. The conversation is dominated by speed, scale, and technical improvement. What is missing is a parallel conversation about authority. Who defines accuracy in a language that is not theirs? Who determines what counts as acceptable interpretation? Who has the power to intervene when the system produces something that is technically functional but linguistically or culturally wrong?

These are not edge questions. They are foundational. Sign languages are not simplified versions of spoken language, and they are not uniform across communities. They carry structure, variation, and meaning that cannot be flattened without consequence. When organizations build systems that interpret or reproduce sign language without embedding Deaf professionals into decision-making roles, they are not just building incomplete products. They are establishing governance models that exclude the very expertise they depend on.

This becomes an operational issue the moment these systems intersect with the workforce. HR functions are often where these decisions surface first, even when they were not made there.

They are establishing governance models that exclude the very expertise they depend on.

Screening tools begin to filter candidates based on communication patterns that assume spoken language norms. Training systems are deployed with accessibility treated as a feature instead of a baseline requirement. Performance metrics reward behaviors that rely on auditory or real-time verbal interaction. None of these choices are typically labeled as exclusionary. They are framed as efficiency. But efficiency measured against the wrong baseline produces predictable outcomes.

The problem is not that organizations intend to exclude. The problem is that they treat accessibility as something that can be layered onto a system rather than something that must shape the system from the start. By the time accessibility is considered, the architecture has already limited what is possible. At that point, Deaf professionals are brought in to validate, adjust, or improve, but not to define. That distinction matters because it determines who holds authority over the system's behavior.

This is why the work emerging from **SLxAI** signals a shift that organizations should pay attention to. The focus is not limited to improving outputs or refining models. It is centered on governance. That includes defining standards, establishing accountability, and determining who has decision-making power over how sign language is represented and used in AI systems. This approach reframes accessibility from a downstream concern to a structural one.

Organizations that ignore this distinction will continue to move quickly, but they will accumulate risk in ways that are harder to measure. Systems that appear to function well can still produce long-term harm if the underlying assumptions are flawed. When those systems are tied to hiring, training, or communication, the impact extends beyond usability

into credibility and trust. Rebuilding that trust after deployment is significantly more difficult than designing for it at the outset.

THE PATTERN REPEATS

As AI systems expand into domains that rely on specialized human expertise, the same pattern will repeat.

The broader implication is not limited to sign language AI. As AI systems expand into domains that rely on specialized human expertise, the same pattern will repeat. If the people who hold that expertise are not embedded into governance structures, the system will reflect only a partial understanding of the domain it operates in. That gap will not always be visible immediately, but it will surface over time through inconsistencies, misalignment, and resistance from the communities affected.

The question organizations need to ask is not whether their systems are improving. It is whether the people whose knowledge those systems depend on have meaningful authority in how they are built and deployed. Without that, accessibility is not being solved. It is being approximated under conditions that remain outside the control of the people most impacted.

That is not a technical limitation. It is a governance choice.