This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/10/29/more-compute-same-question-rylo-nvidia-and-whose-signing-gets-scaled/.

More Compute, Same Question: Rylo, NVIDIA, and Whose Signing Gets Scaled

Heather Grizzle · October 29, 2026

Rylo will train sign language AI on NVIDIA supercomputers. Its terms say Rylo Sign input may train its models. Relay calls are protected. Signers are not.

İçindekiler

1. The rule behind every relay call
2. From a side project to a $500 million relay company
3. What NVIDIA changes, and what it does not
4. One sentence in Annex 2
5. What the privacy policy does not say
6. The dictionary is the better example, and Rylo already built it
7. Open source, noncommercial, and now inside a company
8. The evidence gap has not moved
9. Involvement is not authority
10. What a buyer should require
11. Questions put to Rylo
12. The tokenizer does not ask

On 21 October 2026, Rylo, the company that used to be Nagish, announced it will train its sign language models on NVIDIA's supercomputers. The coverage tells you how fast the models will learn. It does not tell you whose signing they will learn from. The answer to that is sitting in Rylo's own terms of service, in one sentence that most Deaf users will never read.

KISACA

- Rylo (formerly Nagish) holds full FCC certification to run captioned telephone and IP Relay calls with no human call assistants, earns per-minute payments from the federal relay fund, raised $85 million in June 2026 at an estimated $500 million valuation, and bought the sign language translation company sign.mt in October 2025.
- On 21 October 2026, Calcalist reported that Rylo will train its sign-to-text models on NVIDIA's supercomputing infrastructure using NVIDIA's Cosmos Tokenizer. None of the coverage says what signing the models are trained on, who the signers are, or what they agreed to.
- Rylo's terms of service, updated 8 October 2026, say that what a user puts into Rylo Sign "may be used to train or improve our models." Federal relay calls are carved out of that clause. Sign language input is not.
- Rylo's privacy policy, updated 4 October 2026, contains no clause on video, faces, biometric data or sign language, and offers no opt-out from training.
- Rylo's public sign language claims, including "the most accurate ASL dictionary available online," rest on no published independent evaluation. NCG found no independent benchmark for Rylo Sign in August 2026 and has found none since.

RYLO, BY THE NUMBERS

$85M
raised in June 2026, at an estimated $500M valuation

0

1

human call assistants on its federally funded relay calls

sentence in the terms that opens Rylo Sign input to model training

0
privacy policy clauses on video, faces or sign language

0
published independent evaluations of Rylo Sign

7
written questions put to Rylo

Source: FCC DA 26-566 (9 June 2026); Business Wire (4 June 2026); Calcalist (9 June and 21 October 2026); Rylo Terms of Service (8 October 2026) and Privacy Policy (4 October 2026). Checked 29 October 2026.

The rule behind every relay call

I have used relay services for years, and there is a promise underneath every one of those calls that most hearing people have never heard of. Federal rules say a relay provider may not disclose what was said on a relayed call, and may not keep a record of it once the call is over. It does not matter what the call was about. A call to my doctor, my bank, my kid's school or my lawyer ends, and the content is supposed to end with it. That rule is the reason Deaf people trust relay at all. It is written into the Code of Federal Regulations at 47 CFR 64.604, and it is enforced by the Federal Communications Commission.

I want to say up front that I know some of the people who built Nagish. I have watched the company grow from a captioning app with a good idea into one of the most interesting businesses in our space, and I respect what it has done for hard of

hearing people who were tired of waiting on a human captioner to catch up. That is part of why this piece exists. A company with this much money, this much reach and this much goodwill in the Deaf community sets the norm for everyone who comes after it. If the norm is wrong, it is better to say so while it is still being set.

Here is the norm I am worried about. Rylo's relay calls carry the old promise, because federal law makes them. Rylo's sign language product does not carry it, because nothing makes it. And the sign language product is the one Rylo just announced it is scaling.

From a side project to a $500 million relay company

Nagish started, by its chief executive Tomer Aharoni's account to Calcalist, as a side project in 2019. Aharoni and his co-founder Alon Ezer built an app that captioned phone calls in real time and turned a user's typed replies into speech, so a Deaf or hard of hearing person could make a call without routing it through a call center. That design choice is the whole business. Under the FCC's Internet Protocol Captioned Telephone Service program, an approved provider is paid out of the federal Telecommunications Relay Services Fund for every eligible minute it captions, Calcalist reports the company is paid per minute of conversation, at rates regulators reset every three years, and that it became profitable shortly after its first revenues.

The company moved fast through the regulatory gate. Communications Daily reported conditional FCC certifications for Nagish's captioned telephone service in January 2024. On 9 June 2026, the FCC's Consumer and Governmental Affairs Bureau granted Nagish full certification to provide both captioned telephone service and IP Relay "on a fully automated basis," noting that the company "currently does not employ communication assistants and does not operate any call centers." Five days earlier the company had announced it was now called Rylo and had raised $85 million, from General Catalyst's Customer Value Fund and Canaan, with existing investors joining. Calcalist put the valuation at about $500 million and reported a target of $100 million in annual revenue by the end of 2026 and $1 billion by 2028, from a company of 32 employees.

NAGİSH TO RYLO

2019
A side project

Tomer Aharoni and Alon Ezer begin building a real-time call captioning app, by Aharoni's account to Calcalist.

January 2024

Conditional FCC certification

Nagish receives conditional certification to provide IP captioned telephone service, paid from the federal relay fund.

22 October 2025

Buys sign.mt

Nagish acquires sign.mt, a sign language translation company. Its founder, Dr. Amit Moryossef, becomes Head of Research.

4 June 2026

Becomes Rylo, raises $85M

The rebrand announces Rylo Sign, an AI platform for translation between "dozens of signed and spoken languages."

9 June 2026

Full FCC certification

Fully automated captioned telephone and IP Relay, with no communication assistants and no call centers.

4 and 8 October 2026

New privacy policy and terms

The terms state that Rylo Sign input may be used to train Rylo's models.

21 October 2026

NVIDIA partnership

Rylo will train sign language models on NVIDIA supercomputing infrastructure and launches Rylo Desktop in private beta.

None of that is a criticism. A Deaf-serving company that figured out how to make the relay fund pay for automation, and then made itself profitable doing it, has done something hard. It also means Rylo's sign language work is not a startup gamble. It is funded research sitting on top of a regulated revenue stream, run by a company that now has the incentive and the resources to build sign language models at a scale few in this field can match.

What NVIDIA changes, and what it does not

Calcalist's 21 October report describes the partnership in technical terms. Rylo's research team will train its sign language recognition models, the ones that read a person's signing and turn it into text or speech, on NVIDIA's supercomputing systems. The team is using NVIDIA's Cosmos Tokenizer, which compresses video into a compact representation that GPT-style language models can work with, and NVIDIA's NeMo Megatron Bridge software. Moryossef, who founded sign.mt before Rylo bought it, told Calcalist: "Access to NVIDIA's compute gives us the ability to train these models at greater scale and precision." The same report covered the launch of Rylo Desktop, a free Mac app that captions Zoom, Google Meet and Microsoft Teams meetings, in private beta in the United States.

Read that again with a Deaf signer's eyes. A tokenizer is a machine for turning video of a person into training material. Compute is the thing that lets you use more of it. What the announcement changes is scale. What it does not change, and does not even mention, is where the video comes from. Not one of the reports I found, from Calcalist, Techwire Israel or the outlets that repeated them, names a dataset, a sign language, a group of signers, a consent process or a payment. The scale is announced. The supply is not.

That silence is not unusual in this industry. I made the general case in "The Corpus Question" last month: every sign language model is built out of real people's hands, faces and bodies, and the questions that matter are whose signing it is, what they agreed to, what they were paid, and whether they can take it back. What is unusual about Rylo is that the answer to the first question is partly written down, and it is not in a press release. It is in the terms of service.

One sentence in Annex 2

Rylo's terms of service were last updated on 8 October 2026. Near the end is Annex 2, the Rylo AI Supplementary Terms. It defines Rylo Sign as "Rylo's AI-powered sign language translation service." Then, under a heading about improving Rylo AI, it says that the input a user provides to Rylo Live, Rylo Live Transcribe or Rylo Sign, and the user's usage of those services, "may be used to train or improve our models and Rylo AI, consistent with applicable law."

The same paragraph carves out federal relay. The content of calls processed through the Telecommunications Relay Services program will not be used or retained by Rylo or its partners beyond the duration of the call, except as required by law. That carve-out is not generosity. It is the 47 CFR 64.604 rule I started with, showing up in the contract because the law requires it.

So look at what sits on each side of the line. A phone call I caption through Rylo is protected by federal law, ends when I hang up, and never becomes training data. If I open Rylo Sign and sign into my camera, the input is my face, my hands, my body and whatever I happened to be saying, and the terms say it may be used to train Rylo's models. The difference is not that one is more private than the other. The difference is that one is regulated and the other is not.

A relay call through Rylo	Signing into Rylo Sign
Governed by FCC relay rules, 47 CFR 64.604	Governed by Rylo's own terms and privacy policy
Content may not be kept beyond the call	No retention period stated for sign input
Excluded from model training in Rylo's terms	"May be used to train or improve our models"
Paid for by the federal relay fund	Free to use
Captions are speech, not your body	Input is video of your face, hands and body

There is one more document worth reading next to Annex 2. Rylo publishes a separate disclosure for data it receives from Google accounts, which Google requires of apps that use its APIs. That page says Google user data "is never used to develop, improve, or train generalized AI or machine learning models," and that it is excluded from any training data even where the terms of service would otherwise allow it. That is a clear, strong, enforceable promise. Rylo wrote it because Google demanded it. Nobody with that kind of leverage has demanded it on behalf of Deaf signers.

What the privacy policy does not say

The privacy policy, last updated 4 October 2026, is where a user would expect to find the details. It says Rylo may request access to device features "including your mobile device's contacts, microphone, or other features," and that it "may record and store

live conversations that you make available to the Services." It says Rylo does not retain the audio or transcriptions of connected calls in the ordinary course. It describes creating "anonymous, aggregated, or de-identified data" and states that such data is not personal information.

It does not say anything about video. It does not mention cameras, faces, facial expressions, body movement, biometric data or sign language. It does not offer a way to opt out of model training. On AI features generally, it points to the supplementary terms, and adds: "If you do not wish to share information with the third party providing the AI feature, please do not use such feature."

For a captioning company, a privacy policy built around audio made sense. For a company that is now asking signers to point a camera at themselves, it is out of date. Signed language is not like speech in one important way: you cannot sign anonymously. The grammar is on the face. Eyebrows, mouth shapes, eye gaze and head movement carry meaning, so any system that reads signing has to read a face. A promise to de-identify signing video is a promise to remove the part of the data the model most needs, and a buyer or a user is entitled to ask how that works in practice.

This is also where the law gets less settled than Rylo's terms make it sound. "Consistent with applicable law" carries a lot of weight in a country where some states regulate biometric identifiers and most do not. Illinois's Biometric Information Privacy Act, for example, requires written consent before a company collects a "scan of face geometry." Whether a sign recognition pipeline that tracks facial landmarks is collecting one is a question I have not seen a court answer for signing, and I am not a lawyer. It is the kind of question a company scaling sign recognition on a supercomputer should answer in public before a regulator or a plaintiff answers it for them.

WHERE A SİGNER'S VİDEO CAN GO UNDER RYLO'S CURRENT DOCUMENTS

Sign into the camera
A user opens Rylo Sign and signs. The video captures hands, face and body.

→ Rylo Sign input

Rylo Sign input

Is it relay-call content? Only relay content is barred from training by the terms. Sign input is not relay content, so it takes the no branch.

Yes → Discarded No → Training pool

Training pool
Annex 2: input and usage "may be used to train or improve our models." No retention period, no opt-out and no biometric clause are stated.

→ Tokenized and trained

Tokenized and trained
Per Calcalist, Rylo trains on NVIDIA infrastructure using the Cosmos Tokenizer, which turns video into a compact representation for language models.

Discarded
Relay call content is not used or retained beyond the call, as federal rules require.

To be careful about what this diagram does and does not show: it maps what Rylo's published documents permit. It is not evidence that Rylo has trained on any user's signing, and I have no information that it has. A permission in a contract is not a practice. But a permission is what a user agrees to, and it is what a regulator, a school district or a hospital will be held to when something goes wrong.

The dictionary is the better example, and Rylo already built it

In fairness, Rylo has already shown it knows how to do some of this well. The Rylo Sign Dictionary, an ASL dictionary in beta, is unusually open about where its signs come from. Its about page says signs are chosen and maintained by Deaf sign language experts, who either record themselves or upload existing video. Every video is passed through an AI avatar filter so the original signer cannot be identified, and a second annotator checks the result against the original to make sure the sign is still

legible. The page lists every source it draws on, with its license, and says plainly that the sign videos and their links to concepts are licensed for non-commercial use only.

That is the right instinct: Deaf people choosing the signs, a human check on what the machine did, and a license that keeps the signers' work from being resold. What it does not describe is the consent the signers gave, what they were paid, whether their videos are used to train Rylo's models, and whether they can withdraw. Those are the four questions that matter for the dictionary, and they are the same four questions that matter, with far more people involved, for everyone who signs into Rylo Sign.

Open source, noncommercial, and now inside a company

There is a second inheritance in Rylo Sign that deserves attention. Before the acquisition, sign.mt was a research project with an unusually honest public paper. Moryossef's system description, published at the 2024 Conference on Empirical Methods in Natural Language Processing, called the application "by no means a finished product." It said the choice of models for each module was "primarily opportunistic," driven by what was available rather than by performance metrics or user evaluations. It released the code under a Creative Commons Attribution, NonCommercial, ShareAlike 4.0 license, and reported usage figures, not accuracy figures.

That was a reasonable way to run a research project. It is not a reasonable way to describe a commercial product, and Rylo's marketing has moved on without the caveats. The Rylo Sign page today describes Sign Translate as supporting "most signed and spoken languages, including ASL and BSL," calls the dictionary "the most accurate ASL dictionary available online," and describes Sign Translate as open source. Under a noncommercial license, the community that contributed to and relied on sign.mt's openness is entitled to know what open means now that the code sits inside a company valued at half a billion dollars: which parts remain open, under which license, and whether models trained on NVIDIA's infrastructure will be released on the same terms as the research that made them possible.

The evidence gap has not moved

In August 2026, NCG published "The Evidence That Wasn't And Still Isn't There," a search of the required sources for any independent measurement of accuracy, error

rates or user comprehension for the main sign language AI systems on the market. Rylo Sign was one of them, alongside Signapse, Sorenson, Sign-Speak, SignGemma and others. We found no independent benchmark for any of them.

Two months, one rebrand, a supercomputer partnership and a "most accurate" claim later, I have checked again and found nothing new. The funding announcement says Rylo Sign is "powered by the latest research in the field" and gives no accuracy figure. The NVIDIA coverage gives none. The product page gives none. More compute is a reason to expect better models. It is not evidence that the models are good, and it is not evidence that they work for the people who will be read by them: Black ASL signers, older signers, DeafBlind people signing in a tactile or close-vision space, children, people signing fast, people signing in bad light.

Recognition raises the stakes compared with the avatar products I have written about, too. When an airport's AI signer gets a train wrong, a Deaf passenger can at least see that something is off and check the board. When a recognition model gets my signing wrong, the person who sees the error is the hearing person on the other end, and they have no way to know. I am the one being misquoted, and I am the one who cannot see it happen.

What Rylo has said in public, and what it has not

Question	What Rylo has said	What Rylo has not said
Training data for sign models	Sign input may be used to train or improve its models	Which datasets and which signers the current models use
Consent	Google account data is never used to train models	Any consent process for signers whose video is used in training
Biometrics	Nothing	Whether facial or body data from signing is treated as biometric
Retention and opt-out	Relay content is not kept beyond the call	A retention period or training opt-out for Rylo Sign
Accuracy	Calls its dictionary the most accurate ASL dictionary online	Any independent evaluation of Rylo Sign or Sign Translate
Openness	Sign Translate is open source	Which license now applies and whether new models will be released
Deaf authority	Built for and with the Deaf and hard of hearing community	Which Deaf people can approve or stop a model release

Source: Rylo website, Terms of Service and Privacy Policy; Business Wire; Calcalist; PR Newswire; FCC DA 26-566. Checked 29 October 2026.

Involvement is not authority

Rylo says its products are "built for and with" the Deaf and hard of hearing community, and the sign.mt acquisition announcement promised work "developed with and for the deaf community" under the banner "Nothing about us, without us." Calcalist reported that several of the company's 32 employees are deaf, and the Rylo Sign page lists a team of Deaf and hearing researchers, engineers, linguists and

designers, without saying which members are Deaf. I believe the involvement is real. I know some of these people.

But involvement and authority are different things, and the difference shows up at the moment a model is ready to ship. Who decides that a sign recognition model is accurate enough to put in front of a Deaf person's employer, doctor or child's teacher? Who decides what signing goes into the training pool, and whether a user's camera input is part of it? Whose "no" stops the release? None of Rylo's public materials names a Deaf person with that authority. Until one is named, "with" describes the team, not the governance.

What a buyer should require

Most Deaf people will meet Rylo Sign as individual users, clicking through terms on a phone. But the larger exposure is institutional. A school district, a state agency, a hospital or an employer that adopts a free sign recognition tool, or tells Deaf staff and students to use one, inherits the vendor's terms on behalf of people who never saw them. Before any institution relies on Rylo Sign, or any comparable product, it should get the following in writing.

1. The provider MUST state whether video, pose data or any derivative of a user's signing is used to train or improve its models, and MUST offer an opt-out that is on by default for institutional accounts.
2. The provider MUST disclose the datasets used to train each model in the deployed product, the license each dataset was collected under, and whether the signers in it consented to commercial use.
3. The provider MUST state a retention period for sign input and for any derived data, and MUST delete both on request.
4. The provider MUST state whether it treats facial and body data captured from signing as biometric data, and how it complies with state biometric privacy laws where users live.
5. The provider MUST publish independent accuracy evidence for the specific languages, signer populations and settings the institution will use the product in, with error rates reported separately for Deaf signers who are Black, older, DeafBlind or children where those groups are among the users.

6. The provider MUST NOT present recognition output as a record of what a Deaf person said in any setting with legal, medical, educational or employment consequences unless the Deaf person can review and correct it.
7. The provider SHOULD name the Deaf people who hold authority to approve or block a model release, and what that authority covers.

None of these is exotic. Several of them are things Rylo already does for someone else: it excludes relay content from training because the FCC requires it, and it excludes Google data because Google requires it. The ask is to extend the same treatment to signers, without waiting to be required.

Questions put to Rylo

On 25 October 2026, before this piece was written, NCG put seven questions to Rylo in writing, addressed to Tomer Aharoni and Dr. Amit Moryossef, ahead of a full SLAT governance record on Rylo Sign of the kind NCG has published on Signapse and Sign-Speak. Rylo has until 5:00 PM Eastern on 5 November 2026 to answer. Its response will appear with that record in full, and any section of this piece the answers bear on will be revised. In short, the questions are these.

1. Which Rylo Sign models and directions are available to the public today, in which sign languages and in what output form, and is Rylo Sign deployed or piloted by any transit operator, broadcaster, school, university or public body? Which uses, if any, does Rylo exclude?
2. Which datasets trained the Rylo Sign models and supplied the Sign Dictionary's videos, whose signing do they contain, and what did those signers consent to and receive? Can a signer withdraw?
3. Does Rylo Sign capture camera video of users, and if so is it used for training under Annex 2, how long is it kept and which providers receive it? Will Rylo extend to Rylo Sign the no-use, no-retention protection it applies to relay calls?
4. Has any Rylo Sign output been evaluated for intelligibility or accuracy with Deaf raters, or by anyone independent of Rylo, and what is the basis for describing the Sign Dictionary as the most accurate ASL dictionary available online?

5. Who are the Deaf researchers on the team, can a Deaf role halt a release, and will institutional deployments include a defined route to a qualified human interpreter?
6. What route does a Deaf user have, in ASL, to report a wrong rendering, and who answers for an error that reaches the public through an institution's screens or announcements?
7. Which documents govern an institutional deployment of Rylo Sign, will Rylo vary the arbitration, class action waiver and liability cap for public bodies, and does Rylo intend to seek relay fund or other federal funding for any sign language service?

The tokenizer does not ask

I keep coming back to the Cosmos Tokenizer, because it is the most honest part of the announcement. It does what it says. It takes video of a person and turns it into something a machine can learn from, efficiently and at scale. It does not ask whose video it is. It does not know that the eyebrows it is compressing are grammar, or that the hands belong to someone who never read Annex 2. Asking is not the tokenizer's job. It is the company's.

Rylo has the money, the talent, the regulatory track record and, for now, the trust of a lot of Deaf and hard of hearing people. It has also already written the kind of promise I am asking for twice, once because the FCC made it and once because Google did. The third one, for signers, is the one that would mean the most, because nobody is making Rylo write it. If the company that knows our community best will not put that promise in its terms before it scales, I am not sure who will.

METHOD AND DİSCLOSURE

Novara Consulting Group is independent of Rylo, of NVIDIA and of every other company named here, and holds no financial interest in any sign language AI vendor. The author knows some of the people who built Nagish, now Rylo. This piece compares statements Rylo and others have made in public sources, all checked on 29 October 2026, including Rylo′s Terms of Service (last updated 8 October 2026), Privacy Policy (last updated 4 October 2026) and Google User Data Disclosure. It describes what those documents permit, not what Rylo has done in practice; NCG has no information that Rylo has trained models on any user′s signing. Nothing here is legal advice. NCG put seven written questions to Rylo on 25 October 2026, with a response window to 5 November 2026; Rylo′s answers will be published with the full governance record and reflected here.

Nasıl alıntılanır

Grizzle, H. M. (2026, October 29). More Compute, Same Question: Rylo, NVIDIA, and Whose Signing Gets Scaled. Novara Consulting Group. https://www.novaracg.com/tr/2026/10/29/more-compute-same-question-rylo-nvidia-and-whose-signing-gets-scaled/