

Deepgram

NOVA OVERVIEW

Setting a new standard in speech recognition performance.

Traditional Automatic Speech Recognition (ASR) services rely on technologies that are over 50 years old. But in today's hypercompetitive environment, customers need more. More accuracy. More speed. More value. Using the industry's most powerful end-to-end deep learning architecture, Deepgram's award-winning speech recognition and audio intelligence platform is tailor-made for developers building the next generation of intelligent voice applications. Whether you're building a conversational AI voicebot or transforming call center customer experience, our API is 22% more accurate, 23x faster, and 3-7x more affordable than the competition.

Deepgram: Unmatched accuracy. Blazing fast. Enterprise scale.

Nova: The world's most powerful speech recognition API for developers

Our next-gen speech-to-text model, Nova, surpasses all competitors in speed, accuracy, and cost. Real-world testing shows Nova delivers a 22% reduction in word error rate (WER) vs. the nearest provider with inference times that are 23-78x faster than alternative services. But this performance boost doesn't come with increased expense. In fact, Nova is more than 3x more affordable than the closest competitor. These major advances result from years of dedicated research, large-scale distributed training, and in-house data labeling expertise. And it doesn't stop there. Deepgram provides a broad suite of advanced features like domain-specific models, built-in diarization, summarization, and more to enable developers to build with confidence and ship faster.

↑ Highest Performance Speech to Text

⚡ Fastest Inference Time in Industry

💰 Lower Cost

↗ Advanced Speech Understanding

⚙ Custom Model Training

📍 On-prem or Cloud Deployment

9.5% WER Accuracy
(22% Better than nearest)

22% more accurate
than nearest provider

12.1 sec
to process 1 hour of audio

\$0.0043
per minute of audio

3-7x more affordable

A Better Whisper API Than OpenAI's

We also have a fully managed Whisper API that addresses many of the constraints that OpenAI's open source API currently has. Ours supports all five open-source models (not just Large-v2) and is faster, more reliable, and cheaper than OpenAI's. And it also includes built-in diarization, word-level timestamps, and has an 80x higher file size limit. Deepgram's Whisper Large model is 3x faster with >7% fewer word errors than OpenAI's Whisper Large API based on benchmarks with real-world data while also costing 20% less.

Deepgram provides everything developers need to build with confidence and scale from "Hello World" to production fast. [Sign up](#) for a free trial of our Console product, which includes \$200 in credit (around 45K minutes), to start building today.

Build with Confidence:

-  [Console and Free API key](#)
-  [SDK and Tools](#)
-  [Quickstart guides](#)
-  [API reference](#)
-  [Developer use cases](#)

Our Mission: Become the Complete Speech Company

At Deepgram, we believe that speech is the hidden treasure within enterprise data, waiting to be discovered. Our mission is to make world-class language AI not just a possibility but a reality for every developer through a simple API call. We firmly believe that language is the key to unlocking AI's full potential, shaping a future where natural language is the backbone of human-computer interaction. As pioneers in AI-powered communication, Deepgram is committed to transforming how we connect with technology and each other.

Are you looking for an edge over the competition? [Try us for free](#) today and see what all the excitement is about!

2015

Founded

200+

customers

2 trillion

words transcribed

200+

employees worldwide
counting

\$85M+

Raised

For more information visit deepgram.com or [contact us](#) to learn more.

Deepgram

Essential Building Blocks for Language AI