

Lexis

The Omilia multilingual CX AI voice synthesis model that sounds human, responds in real time, and never leaves your platform.

Built for live agentic voice applications in many languages, with no dead air

A voice that sounds great is only half the story, in a live interaction, it has to **perform**. Real-time performance is what separates a conversation that flows from one that stalls.

< 45ms¹

TIME TO FIRST AUDIO

Lexis streams audio as it generates, so there's no dead air, and turn-taking feels natural. The ultra-low-latency pipeline supports barge-in, the agent stops the moment a caller speaks, and holds quality under high concurrency, so automation never feels mechanical or robotic.

25+

LOCALES

Natural, expressive speech across more than 25 locales spanning major world languages, English, Spanish, Portuguese, French, German, Italian, and more, with cloning and streaming in every one. Each locale is tuned for native pronunciation and prosody, while Lexis replicates tone, accent, and speaking style per voice.

Lexis is Omilia's generative text-to-speech model, built to unlock the most authentic Customer Service experiences.

Speech with the expressiveness, warmth, timing, and emotion of a skilled human agent, generated in real time, entirely inside your own platform and adapting to the context and the stakes of every conversation. What makes this possible is that **Lexis is native to Omilia's own, vertically integrated stack**, not assembled from third-party APIs.

01

Authentic experiences that keep customers engaged

Ultra-expressive, context-aware voices streaming in real time that hold the customer in the conversation instead of sending them to an agent.

02

Branded voice identity, equity & control

A voice as recognizable as your logo across all digital channels, made yours, and controlled down to the word.

03

Instant turn-up in minutes

A new branded voice goes live in minutes with a short clip or in seconds with a pre-tuned voice, not months of studio production.

04

Enterprise-grade assurance and predictability for CC

Dependable at the scale of the largest operations, sovereign, compliant, and free of third-party surprises.

Lexis capabilities at a glance, behind the business outcomes

-  **Sub-45ms¹ first audio experiences**
Real-time, ultra-low-latency streaming, for engaging, organically flowing, and perfectly timed conversations that bring agentic and conversational AI customer service agents to life.
-  **Custom voice cloning on the spot**
On-brand voices from a short clip, time-to-launch your custom agent drops from months to minutes, and you can bring your own preferred sound.
-  **25+ locales**
Multilingual coverage across major enterprise markets.
-  **Human-like expressiveness you control**
Generative deep neural networks and NLP for naturally emotive, tone-and-pace-adapting, context-aware delivery, tunable down to pronunciation, pace, and style via lexica, SSML, and voice model parameters.
-  **Pre-tuned CX ambassador personas library**
Ready-to-go, domain-fluent voices with vocal characteristics tuned to the role that fits the stakes.
-  **Zero data leakage**
Enterprise-grade compliance and sovereignty, native to OCP; your voice data never leaves the platform.

In Customer Service, voice quality is not cosmetic. It moves the metrics that decide CX outcomes, and a robotic one is expensive.

A voice customers accept keeps them contained instead of bailing to an agent, protecting the single largest line of contact-center cost.

52%

CUSTOMER SATISFACTION

of consumers stopped buying from a brand after a bad experience.² The voice is the first impression; a natural one removes the friction that a bad one creates.

≈ 1:1

FIRST-CONTACT RESOLUTION

FCR and CSAT move almost in lockstep.³ Clear, well-paced, correctly pronounced speech reduces repeat requests, so every point of resolution is a point of satisfaction, instead of live agent transfer.

The first chance to contain, and to de-escalate

Customers decide how they feel in the first few seconds, and a flat voice is what makes them ask for a human or hang up.

HOW IT WORKS

Prosody shifts with the stakes of the conversation, not just the script.



THE MECHANISM

Generative deep neural networks and NLP read context, tone, and emotion on each word and sentence, then adjust pacing, intonation, inflection, and emphasis, **measured and slower on a sensitive case, brisk and efficient on a routine one.**



THE EFFECT

A calmer, more empathetic delivery defuses tension and reduces the reflex that sends callers straight to an agent, so **more conversations stay on a path to resolution.**

Your voice is a brand asset

BRANDED VOICE IDENTITY

Unmistakably yours, as recognizable as your logo.

A voice is never one-size-fits-all. The right voice depends on who you are, your industry, and the role it plays in the moment; a measured advisor on a banking call sounds nothing like a brisk operator confirming an order. Lexis makes that voice unmistakably yours and keeps it consistent everywhere customers hear you: IVR, outbound, notifications, app, and kiosk.

1

MAKE IT YOURS

Bring or build the voice

Clone a custom, on-brand voice from a short, high-quality reference recording. A few seconds of clean, single-speaker audio is enough to build a usable voice; quality scales with the clip. The same path works whether you supply a voice you already use or commission a new one, there is no multi-week production cycle, only best practices from Omilia's vast field experience. Already use a preferred brand voice, or have a specific sound in mind? Lexis brings it to life instantly.

SO YOU GET a custom-branded voice consistent across every channel, live in minutes.

2

OUT OF THE BOX

Pre-tuned to your industry and roles

Prefer not to start from a recording? Choose a ready-to-go CX Ambassador Persona from the Omilia library. Production-ready, pre-tuned to your industry, and fluent in its terminology, each with a vocal demeanor matched to its role:

The Trusted Advisor

Measured and reassuring, for sensitive, high-stakes conversations.

The Bright Concierge

Warm and welcoming, for greetings, guidance, and hospitality.

The Quick Operator

Crisp and efficient, for high-volume, transactional tasks. And more...

SO YOU GET a fit-for-purpose voice with zero production lead time, just select it and go.

**CONTROL****Two layers of tuning, down to the word**

Two independent layers shape delivery. At the **application layer**, custom lexicons (dictionaries) define exactly how your product catalog, fund and ETF names, or any custom terminology is spoken; SSML shapes rendering term by term. At the **model layer**, voice parameters adjust warmth, pitch, pace, and emotional range. Measured by ASR-over-TTS, output transcribes back at < 3% word error rate (WER).

SO YOU GET exact control over how every term is spoken, no mispronounced names in front of a customer.

Sovereign by design, compliant by default for even the largest of operations

From banks to governments to any SME, where the scale and grade Lexis runs at is as important as how it sounds.

Walled-garden architecture

Native to OCP, with no audio routed to outside speech, language, or voice services and no third-party data sharing, so there is no external surface to audit or explain.

Best-in-class availability at peak

Built to support peak contact-center call concurrency at predictable, best-in-class performance, voice excellence holds even during the highest spikes.

Guaranteed availability & support

Omilia-owned voices, not an experimental library that can change or retire without notice.

Deployment choice

Fully managed SaaS or on-premise for data-residency mandates.

PCI-DSS

SOC 2 Type II

ISO 27001

HIPAA

GDPR

CCPA

Why Omilia Lexis

The **native stack is the difference**, and where third-party TTS chains fall short.

	Third-party TTS API chain	Lexis, native to OCP
Voice quality	Generic, often flat synthetic delivery.	Hyper-realistic, expressive, context-adaptive.
Latency	External round-trips add variable delay.	Sub-45ms first audio chunk inference and as low as 180ms including NW and app latency; streaming, predictable.
Control	Limited control over terms and delivery.	Lexicons, SSML, and model-layer parameters tuning.
Data & sovereignty	Audio routed to external providers.	Stays inside the OCP perimeter; enterprise-compliant.
Brand voice	Generic shared voices; cloning often limited.	Instant cloning or your own custom voice, plus a persona library.
Availability	Experimental libraries; voices can change or retire.	Omilia-owned, SLA-backed, no surprise end-of-life or peak call volume outage.
Cost & deployment	Per-use metering; cloud-only, vendor-dependent.	No third-party metering; managed SaaS or private cloud.

Hear Lexis in your brand voice.

To hear Lexis in your brand voice, contact your Omilia representative.



¹Internal figures. <45ms Excluding application and network latency at TTS inference time only for an average voice agent prompt length. ²PwC, 2025 Customer Experience Survey. ³SQM Group, 2025 FCR/CSAT benchmark.

About Omilia

Omilia is the global standard for AI-driven customer service transformation. Our native Self-Learning Agentic CX platform revolutionizes how enterprises engage with customers, automating interactions with precision, empowering agents in real time, and delivering seamless, personalized experiences across all channels. Powered by deep expertise in developing proprietary Agentic AI technology and multi-layered anti-fraud capabilities, we enable enterprises to move decisively and safely into the era of AI-first contact centers. Omilia's Self-Learning Agentic CX learns from across the entire customer journey, from self-service to live agent interactions, unlocking continuous improvement and breaking the "glass ceiling" of containment that legacy siloed models can't achieve. Omilia is trusted by the world's most demanding enterprises across all industries. Built on over two decades of AI innovation, Omilia delivers measurable outcomes: lower costs, higher efficiency, and unmatched customer satisfaction, all while preserving the human touch where it matters most.