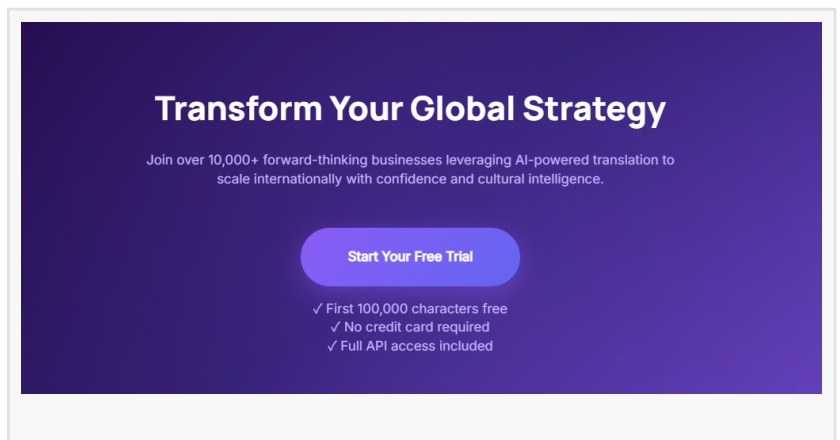


AI.cc Releases AI Translator API, Targeting Enterprises Replacing Legacy Translation Infrastructure

SINGAPORE, SINGAPORE, SINGAPORE, May 21, 2026 /EINPresswire.com/ -- New API delivers neural machine translation powered by multi-model AI routing across GPT-5.5, DeepSeek V4, and Gemini 3.1 Pro, with enterprises reporting 73% cost reduction versus Google Translate API and DeepL API at equivalent quality levels



Transform Your Global Strategy

Join over 10,000+ forward-thinking businesses leveraging AI-powered translation to scale internationally with confidence and cultural intelligence.

[Start Your Free Trial](#)

- ✓ First 100,000 characters free
- ✓ No credit card required
- ✓ Full API access included

SINGAPORE, May 14, 2026 — AI.cc, the Singapore-based unified AI API aggregation platform, today announced the general availability of its [AI Translator API](#), a multi-model neural machine translation service supporting 135 languages designed for enterprises replacing legacy translation infrastructure with AI-native alternatives. The API is available immediately at [api.ai.cc/translate](#), with free-tier access for evaluation and pay-as-you-go pricing for production deployments.

The launch addresses a market inflection point that has been building throughout 2025 and into 2026: the performance gap between dedicated large language models fine-tuned for translation tasks and legacy neural machine translation engines — the technology underlying Google Translate API and DeepL API — has widened to the point where enterprise translation buyers are actively re-evaluating their infrastructure. AI.cc's Translator API delivers LLM-quality translation through a purpose-built endpoint that abstracts the complexity of multi-model translation orchestration behind a single, simple integration.

"Legacy translation APIs were built for a world where translation meant pattern-matching against bilingual corpora," said an AI.cc spokesperson. "LLMs understand context, tone, domain vocabulary, and cultural nuance in a way that fundamentally changes what machine translation can deliver. Our Translator API makes that capability accessible through the same simple integration model enterprises already use for their other AI workloads."

Why Enterprises Are Replacing Legacy Translation Infrastructure

The translation API market has been dominated by Google Translate API and DeepL API for the better part of a decade. Both products deliver adequate quality for general-purpose translation at scale. In 2026, adequacy is no longer sufficient for the enterprise use cases driving the majority of translation API consumption.

Legal document translation, where mistranslated terminology carries material liability risk, requires domain-aware translation that legacy engines handle inconsistently. Financial report localization, where numerical formatting, regulatory terminology, and market-specific conventions vary significantly across geographies, produces errors in legacy translation that require expensive human post-editing. Marketing content localization, where tone, cultural resonance, and brand voice must survive the translation process, is a category where Google Translate and DeepL produce output that brand teams routinely reject as unusable without substantial revision.

Al.cc's platform data shows that enterprises using LLM-based translation for these high-stakes categories reduced human post-editing requirements by 71% on average compared to legacy translation API output — a reduction that translates directly into lower total localization cost even before accounting for the API pricing differential.

Beyond quality, enterprises cite three additional drivers of migration away from legacy translation infrastructure.

Context window limitations in legacy APIs restrict the ability to maintain consistency across long documents. A 50-page contract translated paragraph-by-paragraph loses terminological consistency that a model processing the full document context maintains. Al.cc's Translator API supports document-level context up to 128,000 tokens — sufficient for all but the longest enterprise documents.

Domain customization is unavailable in legacy translation APIs. A pharmaceutical company's translation requirements differ fundamentally from a financial services firm's — specialized vocabulary, regulatory terminology, and formatting conventions are not learnable by legacy engines. Al.cc's Translator API supports system prompt customization that allows enterprises to inject domain context, glossary enforcement, tone specifications, and formatting requirements into every translation request.

Pricing trajectory has shifted against legacy providers. Google Translate API charges \$20 per million characters. DeepL API Pro charges approximately \$25 per million characters. Al.cc's Translator API, routing translation requests through the most cost-efficient LLM capable of meeting quality requirements for each language pair and document type, delivers at \$6–14 per million characters depending on model tier and language pair — a 44–76% reduction versus legacy provider pricing at equivalent or superior quality levels.

Technical Architecture: Multi-Model Translation Routing

The AI.cc Translator API's core technical innovation is its multi-model routing layer, which selects the optimal underlying model for each translation request based on language pair, document type, quality requirements, and cost constraints.

For high-stakes translation requiring maximum quality — legal, financial, medical, and regulatory documents — the API routes to frontier reasoning models including Claude Opus 4.7 and GPT-5.5, which deliver the deepest contextual understanding and most reliable domain-specific terminology handling. For standard business content translation where quality requirements are high but not at the legal-document tier — marketing materials, product descriptions, customer communications — routing targets Claude Sonnet 4.6 and Gemini 3.1 Flash, which provide strong translation quality at mid-tier pricing.

For high-volume, cost-sensitive translation workloads — content moderation at scale, user-generated content translation, real-time chat translation, and batch processing of standard documents — the API routes to DeepSeek V4-Flash and Qwen 3.6-Plus, which deliver translation quality measurably superior to Google Translate and DeepL for most language pairs at \$0.14–0.28 per million input tokens.

Chinese language translation — Simplified and Traditional Chinese, as well as cross-dialect translation — benefits specifically from routing to Chinese-origin models. DeepSeek V4 and Qwen 3.6-Plus consistently outperform Western-origin models on Chinese-English and Chinese-Asian-language pairs, reflecting the linguistic training data advantages of models built by Chinese AI labs. This routing decision alone produces measurably superior output for the Chinese-language workloads that represent a significant share of Asia-Pacific enterprise translation volume.

The routing layer is transparent to API consumers. A translation request specifying a quality tier — standard, professional, or premium — triggers the appropriate model routing automatically. Enterprises that want explicit model control can override routing with a specific model parameter, maintaining the same flexibility available across AI.cc's full API platform.

Language Coverage and Supported Use Cases

The AI.cc Translator API supports 135 languages at general availability, with coverage spanning all major world languages and a significant depth of coverage for Asian, Middle Eastern, African, and Eastern European languages underserved by legacy translation providers.

Tier 1 — Full professional quality (47 languages): English, Mandarin Chinese, Traditional Chinese, Spanish, French, German, Japanese, Korean, Portuguese, Arabic, Hindi, Russian, Italian, Dutch, Polish, Turkish, Vietnamese, Thai, Indonesian, Malay, and 27 additional major world languages. Full professional and premium tier model routing available. Recommended for all enterprise use cases including legal, financial, and regulatory translation.

Tier 2 — High quality (54 languages): Covers regional languages across Southeast Asia, South Asia, the Middle East, Eastern Europe, and Sub-Saharan Africa. Standard and professional tier routing available. Quality materially superior to legacy translation APIs for most language pairs.

Tier 3 — Standard quality (34 languages): Lower-resource languages where LLM training data is less dense. Standard tier routing. Quality comparable to or exceeding legacy translation APIs for general content; not recommended for high-stakes document translation.

Supported enterprise use cases at general availability include document translation with formatting preservation across PDF, DOCX, and HTML inputs; real-time streaming translation for live communication and customer support applications; batch translation for high-volume content processing; glossary enforcement for domain-specific terminology consistency; and tone and style preservation for brand-voice-sensitive marketing localization.

Integration: Built for Enterprise Developer Teams

The AI.cc Translator API is designed for integration by developer teams already familiar with REST API patterns, with OpenAI-compatible formatting ensuring minimal learning curve for teams already using AI.cc's core platform.

A basic translation request requires three parameters — source text, target language, and quality tier — and returns translated text with confidence scores and detected source language:

```
import anthropic_ai_cc as aicc
```

```
client = aicc.Client(api_key="your-api-key")
```

```
response = client.translate(  
text="The indemnification clause shall survive termination.",  
target_language="zh-CN",  
quality="professional",  
domain="legal",  
glossary={"indemnification": "□□□□"}  
)
```

```
print(response.translation)  
# Output: □□□□□□□□□□□□□□□□  
print(response.confidence)  
# Output: 0.97
```

For document-level translation with formatting preservation, the API accepts file uploads directly

and returns translated documents maintaining original layout, typography, and structural elements. PDF, DOCX, XLSX, PPTX, and HTML formats are supported at general availability.

Streaming translation — returning translated text token-by-token as it is generated rather than waiting for full completion — is available for real-time applications including live customer support chat, simultaneous interpretation assistance, and real-time content moderation. Latency for first-token delivery on standard translation requests averages 380 milliseconds on Singapore-origin requests, within the threshold required for real-time interactive applications.

Enterprise customers requiring custom deployment configurations — including private model routing, on-premises processing for sensitive documents, and custom fine-tuned translation models for specialized domains — can engage AI.cc's enterprise team at www.ai.cc/enterprise-plans.

Pricing: Transparent, Tier-Based Structure

AI.cc's Translator API pricing is structured around three quality tiers, each mapped to a model routing configuration and priced per million characters of source text processed:

Standard tier — Routes to DeepSeek V4-Flash and Qwen 3.6-Plus. Best for high-volume, cost-sensitive translation of general content. \$6 per million characters. Compared to Google Translate API at \$20/M characters: 70% reduction.

Professional tier — Routes to Claude Sonnet 4.6 and Gemini 3.1 Flash. Best for business content, marketing localization, and customer communications requiring quality above standard tier. \$14 per million characters. Compared to DeepL API Pro at \$25/M characters: 44% reduction.

Premium tier — Routes to Claude Opus 4.7 and GPT-5.5. Best for legal, financial, medical, and regulatory documents where translation accuracy has material business or legal consequences. \$38 per million characters.

Free tier access includes 500,000 characters of standard-tier translation monthly — sufficient for ongoing evaluation and low-volume production use. No credit card required for free tier registration at www.ai.cc.

Volume pricing for enterprise customers processing more than 100 million characters monthly is available with committed-volume discounts of 20–40% depending on volume tier and quality mix. Contact enterprise@ai.cc for volume pricing discussions.

Availability and Next Steps

The AI.cc Translator API is available immediately for all registered AI.cc users. Existing customers can access the API using their current API key with no additional configuration. New users can

register at www.ai.cc and access free-tier translation immediately upon registration.

Full API documentation, language coverage tables, integration guides, and code examples in Python, JavaScript, Go, and Java are available at docs.ai.cc/translator.

Enterprise customers requiring SLA-backed production deployment, custom domain glossaries at scale, private model routing, or compliance documentation for regulated industries should contact the enterprise team at www.ai.cc/enterprise-plans.

About AI.cc

AI.cc is a unified AI API aggregation platform headquartered in Singapore, providing developers and enterprises with access to 312 AI models — including GPT-5.5, Claude Opus 4.7, Gemini 3.1 Pro, DeepSeek V4, Llama 4, Qwen 3.6-Plus, and more — through a single OpenAI-compatible API. The platform supports text, image, video, voice, code, embedding, OCR, and translation model categories. Additional offerings include the OpenClaw AI agent framework, enterprise SLA plans, AI application development services, AI Web Scraping API, and GEO-optimized SEO and PR services.

Free API access: www.ai.cc Translator API documentation: docs.ai.cc/translator Enterprise plans: www.ai.cc/enterprise-plans

AICC

AICC

+44 7716 940759

support@ai.cc

This press release can be viewed online at: <https://www.einpresswire.com/article/914012065>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.