

Table of Contents

TECHNICAL PROPOSAL.....1

 AFI-2026-124 — AI-Driven Translation & Language Access Services1

COVER LETTER.....1

EVALUATION SECTION 1: COMPANY QUALIFICATIONS AND EXPERIENCE.....2

EVALUATION SECTION 2: AI TECHNOLOGY CAPABILITY, LANGUAGE COVERAGE & ACCURACY11

EVALUATION SECTION 3: SECURITY, COMPLIANCE & DATA PRIVACY13

EVALUATION SECTION 4: IMPLEMENTATION, ONBOARDING & SUPPORT15

EVALUATION SECTION 5: PRICING.....18

APPENDIX SUBMISSIONS20

TECHNICAL PROPOSAL

AFI-2026-124 — AI-Driven Translation & Language Access Services

Submitted by: Bluejay AI Technologies Inc.
UEI: KTEPUMM1APA1 | **CAGE:** 14CX8 | **EIN:** 93-3816473
Address: 4200 Cheeney Street, Santa Clara, CA 95054
Point of Contact: Joseph Barrientos Liba, CEO/CTO — support@bluejayai.net — (978) 799-7069
Submitted via: Indigo Procurement Portal (<https://indigo.bonfirehub.com>)
Date: May 11, 2026

Workstreams Proposed: 1, 2, 4, and 5

COVER LETTER

The Alliance for Innovation
Attn: Procurement Team
AFI-2026-124 — AI-Driven Translation & Language Access Services

Dear Evaluation Committee,

Bluejay AI Technologies Inc., operating as **Interpret AI**, is pleased to submit this proposal in response to AFI-2026-124. We are an AI-native language services company purpose-built for the government market,

combining advanced speech recognition, LLM-powered translation, and desktop publishing capabilities under one SOC 2 Type II-certified platform.

In fewer than three years of operation we have secured five active government contracts — including a custom AI transcription portal for the California Department of Corrections and Rehabilitation (CDCR), human translation and custom InDesign layout work for the California Civil Rights Department, a custom captioning portal for the California High-Speed Rail Authority, a custom court transcription portal for the North Dakota Department of Transportation, and a statewide translation services vendor award in final negotiation with the State of Oklahoma (Statewide Contract SW1071). Each of these engagements demonstrates a direct capability alignment with one or more of the five Workstreams described in this RFP.

We propose on Workstreams 1, 2, 4, and 5. We believe the cooperative contract model AFI and Edge Public have designed is precisely the right vehicle for municipalities that lack procurement bandwidth to run individual solicitations — and that an AI-first provider like Interpret AI can deliver measurably better cost-per-word, cost-per-minute, and time-to-delivery benchmarks than traditional LSPs while maintaining the human review safeguards that high-stakes government content requires. Beyond our four active contracts, we have active bid submissions or evaluations pending across ten additional jurisdictions including Snohomish County (WA), the City of Los Angeles, the First Judicial District of Pennsylvania, Florida’s statewide post-secondary institution network, and federal agencies including NHTSA and the U.S. Air Force — demonstrating a robust and rapidly expanding government pipeline.

We look forward to the opportunity to serve Participating Public Agencies nationwide.

Respectfully submitted,

Joseph Barrientos Liba

CEO/CTO, Bluejay AI Technologies Inc.

EVALUATION SECTION 1: COMPANY QUALIFICATIONS AND EXPERIENCE

1.1 Company Overview

Legal Name: Bluejay AI Technologies Inc.

DBA / Trade Name: N/A

Entity Type: C Corporation, Delaware (File #2450058), incorporated October 5, 2023

California Registration: Entity #B20250061982 (registered April 5, 2025)

Business Address: 4200 Cheeney Street, Santa Clara, CA 95054

SAM.gov Status: Active (UEI: KTEPUMM1APA1; CAGE: 14CX8; expires August 2, 2026)

Size Status: Small Business; CA DGS Micro Small Business (Cert ID: 2045472, exp. May 31, 2027)

Diverse Business Participation: As a certified CA DGS Micro Small Business, Interpret AI champions diverse supplier participation. When subcontracting is required (e.g., specialized human ASL interpreters or regional dialects), our policy is to aggressively source and partner with minority-owned, woman-owned, and veteran-owned Language Service Providers (LSPs).

Primary SIC: 7338 — Secretarial & Court Reporting Services
Registered NAICS: 541519, 541511, 541512, 541930, 518210

Primary Service Lines Relevant to This RFP:

Service Line	Description
Real-time AI captioning & transcription	Live event SRT/VTT, parole hearing transcription, meeting captioning
Multilingual document translation	Government forms, notices, legal documents, InDesign/DTP layout preservation
AI voice interpretation	Real-time multilingual audio rendering for meetings and events
AI translation platform	SaaS platform with REST API, SSO, role-based access, and audit logging
Video/audio dubbing	AI-powered voice dubbing of video and audio content into target languages — replaces the original audio track with a translated, synthesized voice in the target language while preserving timing and pacing
Emergency alert translation	Sub-5-minute turnaround for public safety notifications

Officers and Key Personnel:

Name	Role	Background
Joseph Barrientos Liba	CEO/CTO	M.S. ECE, Carnegie Mellon; ex-Amazon, TikTok, Airbnb
Suk Hyun (Jake) Kim	COO/CBO	Operations and business development
Tianxing (Olivia) Yang	CMO	Marketing and government outreach
Tanner Bogart	CSO	Security and compliance

Additional Clients (Private / Non-Profit / Institutional):

Beyond our government contracts, our platform is actively used by private and institutional clients that demonstrate real-world scale and diverse use-case coverage:

- **Gulf Coast Community Church** — Live event interpretation contract. Interpret AI delivers real-time AI voice interpretation for weekly services attended by a multilingual congregation. Attendees access interpretation by scanning a QR code displayed on-screen, selecting their preferred language, and receiving live translated audio directly on their personal device — no app download required. The platform handles mid-session language switching and code-switching (speakers who shift between languages within a single utterance), maintaining interpretation continuity without manual intervention. Attendees also have access to an AI chatbot during the live session — allowing them to ask questions about what is being discussed, request clarifications, or retrieve prior statements, all in their chosen language. This is

production deployment of the exact participant-facing access model described in Workstream 1, giving us proven live-event infrastructure that government meeting platforms can leverage directly.

- **University Staff and Students** — Our platform is in active use by staff and students at multiple universities for on-demand translation and interpretation services, demonstrating scalability to high-volume, concurrent-user environments with diverse language needs across academic, administrative, and community-facing content.

These clients complement our government portfolio by showing the platform's operational readiness beyond pilot scale.

Government-Specific Certifications and Cooperative Contract Experience:

- SOC 2 Type II — InterCert, audit period 01 Jun 2024 – 06 Feb 2025; Security, Availability & Confidentiality criteria; report available under NDA upon request
 - CA DGS Micro Small Business (Cert ID: 2045472)
 - NAICS 541930 (Translation and Interpretation) registered in SAM.gov
 - HIPAA BAA available for engagements touching PHI
 - This proposal represents our first cooperative contract submission; all prior contracts were awarded via direct agency solicitation
-

1.2 AI Technology and Platform

Underlying Architecture

Our platform is built on a multi-layer AI stack:

1. **Speech Recognition (ASR) with Live Speaker Diarization:** Our ASR layer is powered by the **Soniox API**, which provides state-of-the-art real-time speech recognition with built-in live speaker diarization — identifying and labeling individual speakers as the meeting progresses, not as a post-processing step. This means captions, transcripts, and translated output are attributed per speaker (Speaker 1, Speaker 2, etc.) from the moment they speak, in real time. Post-meeting transcripts are delivered with speaker-segmented output ready for review; agencies can relabel speaker identifiers to actual names as needed. This is a rare capability in live interpretation platforms and significantly reduces the manual cleanup required to produce a usable formal record.
2. **Translation (LLM-Powered, Provider-Agnostic):** Our platform is built on a provider-agnostic AI connector architecture — we have integrated with Gemini, OpenAI, and other leading LLM providers and can route or switch between them as needed. This means we are not locked into any single vendor: if a provider changes pricing, has an outage, or a better model becomes available, we can adapt without disrupting agency service. Currently we route translation through Gemini using carefully engineered prompts tuned for government and municipal

terminology, tone, and domain context. Any language pair is available immediately with optional human review on the same job request.

3. **Real-Time Interpretation:** For Workstream 1 (live voice interpretation), we combine our ASR layer with a streaming translation pipeline: audio input → ASR transcript (latency <800ms) → Gemini-powered translation → audio synthesis (TTS). The entire round-trip latency from source speech to rendered target audio is under 2.5 seconds for major language pairs.

Audio input is flexible by design. Agencies have two options:

- **Desktop app (downloadable):** Captures system audio directly — no microphone required, no acoustic loopback. Works with any audio source playing on the computer. Includes a popout caption window that can be positioned or cast to a second screen, making captions visible to the room without disrupting the primary display.
- **Web app:** Captures audio via the device microphone. If a meeting is playing through laptop or room speakers, the microphone picks it up naturally — no special setup or integration required.

This means Interpret AI works with virtually any existing meeting setup without requiring API integrations, hardware changes, or IT involvement.

4. **Document Ingestion and Format Support:** Our document translation pipeline combines two approaches depending on document type. For standard structured documents (DOCX, PPTX, HTML, plain text, and more), we parse the underlying file structure — XML, HTML, metadata, markup, and text layers — to extract and replace text while preserving layout. For scanned PDFs and documents with text embedded in images, we route through **Azure AI Document Translation + Azure Document Intelligence** (Microsoft, FedRAMP compliant, U.S. data residency), which performs OCR to extract text from images before translating and reconstructing the document. This covers a broad range of formats with a file size limit of 1 GB:

Category	Supported Formats
Microsoft Office	.doc, .docx, .rtf, .xls, .xlsx, .ppt, .pptx
WPS / Chinese Office	.wps, .et, .dps
PDF & Design	.pdf, .ai, .indd, .idml
Open Document	.odt, .ods, .odp
eBook / Help	.epub, .chm
Images (OCR via Azure)	.jpg, .jpeg, .png, .webp, .svg
Subtitles / Captions	.srt, .ass, .ssa, .vtt
Video / Audio (dubbing)	.mp4, .mov, .mp3, .wav — AI dubbing with TTS voice replacement in target language
Localization / Dev	.json, .resjson, .xml, .html, .go, .yaml, .yml, .php, .plist, .stringsdict, .po, .xlf, .xliff
Data / Text	.csv, .tsv, .txt, .md
Academic	.tex, .arxiv

Category	Supported Formats
Archive	.zip (containing any of the above)

Scanned documents and image-embedded text are handled via Azure AI Document Intelligence OCR, which extracts text from images and scanned pages before translation. Output quality for scanned content depends on scan resolution and image clarity. For PDFs and Office files with standard text layers, translation works as expected. For the most demanding layout work — multi-column government publications, ADA-compliant branded reports — we offer Adobe InDesign (.indd/.idml) DTP services where a human specialist rebuilds the layout in the target language, as we currently do for the CA Civil Rights Department (Agreement #25-0014).

Translation quality is consistently high for well-structured documents. Documents with complex multi-column layouts or heavily nested formatting may have minor formatting irregularities in the output; human review is available on any job upon request.

Video/Audio Dubbing

For video and audio content requiring full language replacement, we offer AI-powered dubbing: the original audio track is transcribed via Soniox ASR, translated via Gemini, and re-synthesized as a new audio track in the target language using ElevenLabs — with timing aligned to the original speaker pacing. Output is delivered as a dubbed video file or standalone audio track. This is directly applicable to government training videos, public service announcements, recorded public meetings, and multilingual accessibility requirements for video content.

Model Provenance and Transparency

Component	Source	Proprietary Layer
ASR + Live Diarization	Soniox API	Real-time speech recognition with live speaker diarization
Translation	Gemini (Google)	Government-tuned prompt engineering; any language on demand
TTS (live interpretation)	OpenAI TTS	Real-time audio synthesis for live interpreted output
TTS (dubbing)	ElevenLabs	Post-production dubbing track generation with high-fidelity voice cloning
DTP engine	Custom (proprietary)	InDesign XML parsing, layout reconstruction

Government and Municipal Terminology Handling

We maintain term bases specific to: California government (CDCR corrections vocabulary, CA Civil Rights terminology, CA transportation terminology), federal procurement (FAR clause translation templates), public health (HIPAA-adjacent vocabulary), and general municipal communications. Term bases are agency-specific and stored separately per client to prevent cross-contamination. Agencies can submit

terminology updates through our self-service portal; changes propagate to active glossaries within one business day.

Quality Assurance Methodology

- **Human Review (per-request):** Agencies can request human review on any individual transcription or translation job at submission time. A qualified human reviewer audits the AI output before delivery. No tier commitment or subscription required — it is a simple per-job option.

Known Limitations

Per AFI’s AI Accuracy Disclosure requirements, we disclose the following:

Item	Notes
ASL / sign language	Not offered as a primary service; open to subcontracting arrangements with qualified ASL providers on a per-agency basis upon request
FedRAMP pre-authorization	Not yet certified; SOC 2 Type II available

- **Low-Resource Spoken Languages (WS1):** Per AFI’s AI Accuracy Disclosure requirements, while Interpret AI natively supports 60+ languages for live interpretation, current industry-wide ASR models do not yet achieve municipal-grade accuracy for live spoken dialects of certain low-resource languages (specifically Somali, Amharic, Tigrinya, Khmer, and Hmong). For these specific languages, we advise agencies to utilize traditional human interpreters until ASR models achieve acceptable Word Error Rates (WER). All of these languages are fully supported for Written Document Translation (WS2).

Known accuracy limitations: ASR accuracy varies by language and acoustic conditions. Our Soniox ASR layer achieves sub-7% WER for most high-demand municipal languages (Spanish, French, Mandarin, Korean, Vietnamese); Arabic is higher at ~16% WER due to dialect diversity and is a known limitation of all current ASR providers. Heavily accented speech, low-quality audio input, and highly technical domain vocabulary (legal citations, medical terminology) may reduce accuracy. Human review is available on every job — we recommend it for high-stakes legal, medical, or formal public-record content.

Platform Updates

Bug fixes are deployed immediately as discovered. Any workflow-level changes are announced to agencies in advance — no surprise changes to how the platform operates.

Federal and State Compliance Certification

Interpret AI formally certifies that our platform and services support Participating Public Agency (PPA) compliance with Executive Order 13166, Title VI of the Civil Rights Act of 1964, the Americans with Disabilities Act (ADA), Section 508 of the Rehabilitation Act, and applicable state language access statutes. All digital interfaces are Section 508 and WCAG 2.1 AA compliant.

1.3 Language Coverage (Summary)

Full detail is provided in Appendix E (Language Coverage Matrix). Summary:

- **Live Voice Interpretation (WS1):** 60 languages supported via Soniox ASR (see full list in Appendix E). Covers all high-demand municipal languages including Spanish, Mandarin, French, Arabic, Vietnamese, Korean, Portuguese, Russian, Hindi, Tagalog, and more.
- **Document Translation (WS2):** Any language, on demand. Human review is available as an optional add-on on any individual job request — agencies do not need to commit to a review tier in advance.

Top-supported languages (Tier 1–2 full coverage across all service types): Spanish, Mandarin (Simplified and Traditional), Cantonese, Vietnamese, Portuguese (Brazilian and European), Arabic (MSA and Levantine), Russian, French, German, Japanese, Korean, Hindi, Tagalog/Filipino, Polish, Ukrainian, Persian/Farsi, Turkish, Bengali, Urdu.

1.4 Government and Municipal Experience

Contract 1 — CDCR / Board of Parole Hearings (California)

Agency: CA Department of Corrections and Rehabilitation

Contract Type: AI-assisted hearing transcription

Term: January 2026 – present (ongoing)

POC: Mariann Isola, mariann.isola@cdcr.ca.gov, (279) 210-3481

Services: Full custom portal for parole hearing transcription — AI generates a draft transcript from pre-recorded hearing audio (with speaker diarization), which human reviewers then review, correct, and approve within the same platform before the transcript is finalized as a legal record

Relevance: Demonstrates end-to-end AI + human review workflow built for a high-stakes government use case; production deployment of a custom transcription portal in a California state agency context

Contract 2 — North Dakota Department of Transportation

Agency: ND DOT, Court Transcription Division

Contract Type: Court transcription services

Term: July 2025 – July 2027 (2-year base)

POC: J. Sprecher (jsprecher@nd.gov), A. Glass (aglass@nd.gov)

Services: Full custom portal for court transcription — AI generates a draft transcript from pre-recorded audio with speaker diarization, which human reviewers review, correct, and approve within the same platform before finalization as a legal record

Relevance: Second active custom transcription portal deployment; demonstrates multi-state delivery capacity and legal-domain accuracy

Contract 3 — California Civil Rights Department

Agency: CA Civil Rights Department

Contract Type: Document translation + DTP/InDesign layout

Agreement #: 25-0014

Value: Up to \$200,000 (multi-year)

Term: October 2025 – September 2028

POC (primary): Rashida Harmon, rashida.harmon@calcivilrights.ca.gov, (510) 227-2143

Contract Analyst: Olumide Ojekunle, olumide.ojekunle@calcivilrights.ca.gov

Services: Human translation of official CRD documents and public-facing materials into 13 languages; custom InDesign layout work to preserve visual design and ADA compliance across all target languages

Relevance: **Direct past performance match for WS2 (Written Document Translation)**; demonstrates professional human translation at scale and InDesign DTP capability for government publishing standards — rare among language services providers

Contract 4 — California High-Speed Rail Authority

Agency: CA High-Speed Rail Authority

Contract Type: Video captioning — SRT/VTT format

Award Date: April 17, 2026

POC: Tina Aldama, tina.aldama@hsr.ca.gov

Services: Full custom portal for video captioning — AI generates SRT and VTT caption files from meeting recordings with speaker diarization, which human reviewers review and approve before delivery

Relevance: Third active custom transcription portal deployment; demonstrates captioning workflow for high-profile government meeting recordings; directly relevant to WS1 (post-meeting multilingual transcript generation)

Contract 5 — State of Oklahoma (Statewide Contract SW1071, IC-2490)

Agency: Oklahoma Office of Management and Enterprise Services (OMES)

Contract Type: Statewide translation services vendor — human translation and interpretation

Status: Award in final negotiation (May 2026)

POC: Tatyana Orrock Nunes, Deputy General Counsel, OMES Legal Division —

Tatyana.OrrockNunes@omes.ok.gov, (405) 522-6236

Services: Human translators and interpreters deliver all services under this contract. Our linguists work through Interpret AI's secure delivery portal, which uses locally-run AI tools as internal productivity aids (terminology lookup, formatting assistance, glossary management) to support the human-produced work product. All deliverables are human-translated and human-verified.

Relevance: Demonstrates Interpret AI's capability to deliver fully human-led translation services at statewide scale, with our portal infrastructure supporting human linguists rather than replacing them — a model directly applicable to agencies that require human-delivered work product within an efficient technology-enabled workflow.

1.5 Geographic Coverage and Scalability

Interpret AI is a cloud-native platform deployed on Vultr U.S. data centers (no offshore data processing). Our delivery model is geographically unrestricted — a municipality in rural Montana procures the same platform access and SLAs as Los Angeles. All services are delivered remotely:

- **Document translation** is delivered via secure portal upload/download or API

- **Live voice interpretation** is delivered via browser-based WebRTC session (no software install required for attendees)

Scalability: Our platform runs on pre-provisioned Vultr infrastructure sized to handle concurrent demand across our client base. We have no geographic or time-zone restrictions on service delivery. Emergency Workstream (WS5) is 24/7/365. Our infrastructure has been validated through institutional deployments serving university staff and student populations with unpredictable demand patterns, providing confidence in our ability to serve municipalities of varying sizes.

Language Availability: Live voice interpretation is available in 60 languages (Soniox-supported) 24/7. Document translation is available in any language via Gemini with no catalog restriction.

1.6 References and Past Performance

Reference 1

Agency: California Civil Rights Department

Contact: Rashida Harmon, Contract Manager

Title: Contract Manager

Phone: (510) 227-2143

Email: rashida.harmon@calcivilrights.ca.gov

Contract Value: Up to \$200,000

Duration: October 2025 – September 2028

Services: Document translation (13 languages), InDesign DTP formatting, ADA-compliant layout

Volume: Ongoing; multiple documents translated monthly across 13 language pairs

Languages: Spanish, Chinese (Simplified), Vietnamese, Korean, Tagalog, Armenian, Farsi, Punjabi, Khmer, Japanese, Russian, Arabic, Hmong

Reference 2

Agency: California Department of Corrections and Rehabilitation / Board of Parole Hearings

Contact: Mariann Isola

Title: Contract POC

Phone: (279) 210-3481

Email: mariann.isola@cdcr.ca.gov

Contract Value: Active (value not publicly disclosed)

Duration: January 2026 – present

Services: Custom transcription portal — AI draft from pre-recorded audio, human review and approval workflow, formatted legal transcripts

Volume: Ongoing; multiple hearings per week

Languages: English (audio input may include multilingual speakers)

Reference 3

Agency: North Dakota Department of Transportation

Contact: J. Sprecher

Title: Contract POC

Phone: Available upon request

Email: jsprecher@nd.gov

Contract Value: Active (value not publicly disclosed)

Duration: July 2025 – July 2027

Services: Court transcription services; verbatim formatted transcripts for legal record

Volume: Ongoing

Languages: English

Note on RFP Requirement (5 contracts): Interpret AI was incorporated in October 2023. We have four active government contracts plus a fifth statewide contract (Oklahoma SW1071) in final negotiation. All five were secured within our first 26 months of operation, span four states, and cover four distinct service categories directly relevant to this RFP (transcription, document translation, video captioning, and human translation services). We believe the depth, relevance, and ongoing status of these engagements is a stronger indicator of capability than a longer list of older or completed contracts from a larger legacy provider. We are actively registered on Procurated and will request agency reviews from each POC upon submission.

EVALUATION SECTION 2: AI TECHNOLOGY CAPABILITY, LANGUAGE COVERAGE & ACCURACY

2.1 Accuracy and Quality Standards

Published Accuracy Rates

Language Pair	Service Type	Accuracy Rate	Measurement Method
---------------	--------------	---------------	--------------------

ASR Accuracy — Soniox Independent Benchmark (March 2025)

Our ASR layer uses Soniox, which published a comprehensive third-party benchmark (March 2025) comparing 10 providers across 60 languages using real-world YouTube audio with human-verified ground truth. Soniox achieved the **lowest Word Error Rate (WER) across most languages** — outperforming OpenAI Whisper, Google, AWS, Azure, NVIDIA, Deepgram, AssemblyAI, Speechmatics, and ElevenLabs. Selected results:

Language	Soniox WER	Next Best Competitor	Competitor WER
Spanish	5.3%	Azure	9.5%
French	5.7%	OpenAI	9.4%
Italian	4.5%	Google	9.6%
Portuguese	4.8%	OpenAI	9.0%
English	6.5%	Azure	9.4%
Mandarin (Chinese)	6.6%	Azure	10.1%
Korean	4.3%	Azure	11.2%

Language	Soniox WER	Next Best Competitor	Competitor WER
Arabic	16.2%	Speechmatics	22.9%
Vietnamese	5.4%	Azure	11.2%
Hindi	7.4%	ElevenLabs	13.0%
Russian	6.2%	Deepgram	8.0%
German	7.0%	ElevenLabs	10.3%

Full benchmark: Soniox Inc., *Soniox Speech-to-Text Benchmarks*, March 2025 — <https://soniox.com/media/SonioxSTTBenchmarks2025.pdf>

Document Translation Quality: Driven by Gemini’s multilingual capabilities. Human review available on any job upon request.

Translation quality is driven by Gemini’s multilingual capabilities. Human review is available on any job upon request.

Quality Assurance Process

Our document translation workflow is straightforward:

1. **Translation:** Gemini generates the translation using government-tuned prompts
2. **Delivery:** Output is delivered in the original file format
3. **Human Review (optional, per request):** If requested at submission, a qualified human reviewer audits the translation before delivery

Escalation: Agencies can flag any delivered translation for human review after receipt; we will turn around a reviewed version per the standard human review SLA.

Real-Time Quality Monitoring: For live interpretation (WS1), we display a real-time confidence indicator (green/yellow/red) to the agency administrator. If confidence drops below threshold (e.g., due to background noise or heavily accented speech), the moderator can trigger a human interpreter connection within 90 seconds.

2.2 Real-Time Interpretation Performance Benchmarks (Workstream 1)

Metric	Commitment
ASR transcription latency	< 800ms from end of utterance
Translation latency (major pairs)	< 1,200ms
End-to-end (speech → rendered target audio)	< 2,500ms (major pairs); < 4,000ms (minor pairs)
Live voice interpretation language support	60 languages (Soniox-supported; see benchmark)
Document translation language support	Unlimited — any language via Gemini
Word Error Rate (WER) — major pairs	4.3–7.4% (Korean, Italian, Portuguese, Spanish, French, English, German, Mandarin — per Soniox March 2025)

Metric	Commitment
	benchmark)
Word Error Rate — lower-resource pairs	8–16% (Arabic, Thai, Malayalam, etc. — per same benchmark; still best-in-class vs. competing providers)
Simultaneous connection capacity	Up to 500 concurrent multilingual streams per event
Code-switching / mid-session language switching	Supported — attendees can switch language at any time; speaker language shifts detected automatically
AI chatbot (live session)	Attendees can interact with an AI chatbot during the live event — ask questions, request clarifications, retrieve prior statements — all in their chosen language
Live speaker diarization	Real-time speaker segmentation (Speaker 1, Speaker 2, etc.) as the meeting progresses — not post-processing; agencies relabel to actual names after the fact
Platform uptime SLA	99.95% (see Section 4.3)
Connectivity recovery time (if session drops)	< 30 seconds (auto-reconnect enabled)
Post-meeting transcript generation	Within 60 minutes of meeting end
AI meeting notes & summary	Automatically generated after each session — structured summary of key discussion points, decisions, motions, votes, and action items; delivered in English and any requested target language simultaneously. Designed to accelerate public meeting minutes drafting for city clerks and administrators.

Post-Meeting Deliverables: Within 60 minutes of meeting end, agencies receive: (1) a full verbatim transcript with speaker diarization, (2) multilingual translated transcript in any requested languages, (3) an AI-generated structured meeting summary covering key discussion points, decisions, motions, votes, and action items — in English and any target language simultaneously. The summary is purpose-built to accelerate public meeting minutes drafting for city clerks and administrators, reducing a multi-hour manual task to a review-and-approve workflow.

Supported Meeting Platforms at Launch: Zoom, Microsoft Teams, Webex, Google Meet. Government-specific integration available for Granicus (WebVTT caption injection) within 90 days of contract execution.

EVALUATION SECTION 3: SECURITY, COMPLIANCE & DATA PRIVACY

3.1 Security Certifications

Certification	Status	Notes
SOC 2 Type II	InterCert — audit period 01 Jun 2024 – 06 Feb 2025; Security, Availability & Confidentiality criteria	Report available under NDA upon request

Certification	Status	Notes
HIPAA BAA	Available	Can execute with any PPA that processes PHI
ISO 27001	In process	Targeting certification by Q4 2026
FedRAMP	Not yet authorized	Hosted on Vultr U.S. data centers; FedRAMP Tailored roadmap for 2027
CJIS	Not yet certified	Will pursue upon receipt of a law enforcement or criminal justice use case request
StateRAMP	Evaluating	Will pursue for CA, TX, NY priority states in Year 2

Note on FedRAMP: Our solutions run on Vultr U.S. data centers. For PPAs requiring FedRAMP Moderate/High pre-authorization, we disclose this limitation honestly and will work with the agency to assess alternatives (e.g., FERPA equivalency, SOC 2 + contractual safeguards).

3.2 Data Handling Practices

What Data Is Collected and Processed: - Document translation: source files and translated output are temporarily stored in encrypted S3 (AES-256 at rest) for up to 30 days post-delivery, then auto-deleted unless the agency enables long-term translation memory archiving (opt-in) - Live voice interpretation: audio streams are processed in real time; no permanent audio recording is retained by default; WebRTC streams are not stored; only the generated transcript (if the agency enables post-meeting transcript) is retained

Data Retention and Deletion: - Default retention: 30 days post-job for documents; 7 days for session metadata - Agency-configurable retention: up to 3 years for translation memory archiving (opt-in only) - Right-to-purge: agencies can request full data deletion for their account at any time; deletion confirmed within 72 hours

AI Model Training: We do not use agency content to train or fine-tune AI models without explicit written consent. Client data is isolated and not accessible to model training pipelines. Our agreement with Google explicitly prohibits using submitted content for Gemini model training.

Data Residency: All data processing and storage is U.S.-based. We use Vultr U.S. data centers as our primary infrastructure. No offshore data processing. Sub-processors (Google Gemini, Soniox) process content in U.S. regions only; this is enforced in our data processing agreements with each vendor.

Sub-Processors:

Sub-Processor	Service	Data Region
Vultr	Infrastructure, compute, storage	U.S. data centers
Soniox	ASR, live speaker diarization	U.S.
Google Gemini	LLM-powered translation (live	U.S.

Sub-Processor	Service	Data Region
	interpretation)	
Microsoft Azure (Document Intelligence + Document Translation)	OCR and document translation for scanned/image PDFs	U.S. (FedRAMP compliant)
OpenAI TTS	Live audio synthesis for real-time interpretation	U.S.

Encryption: AES-256 at rest (S3 server-side encryption); TLS 1.2+ in transit (TLS 1.3 preferred; FIPS 140-2 compatible cipher suites available)

Breach Notification: Within 72 hours of discovery of any breach affecting agency data, we will notify the affected agency and AFI in writing with: the nature of the breach, data types affected, estimated scope, immediate containment actions taken, and remediation timeline.

Penetration Testing: Annual third-party penetration test; most recent test completed Q1 2025. Results summary available to agencies under NDA.

EVALUATION SECTION 4: IMPLEMENTATION, ONBOARDING & SUPPORT

4.1 Implementation Plan

Timeline from Contract Execution to Full Deployment:

Day	Milestone
1–3	Contract execution; agency kickoff call; assign dedicated Customer Success Manager
3–7	Agency onboarding form: collect branding preferences, preferred language pairs, department structure, user list, SSO configuration
7–14	Platform provisioning: agency tenant created; SSO configured (SAML 2.0 / Azure AD / Okta); role-based access configured for departments
14–21	Terminology setup: import agency glossaries; configure municipal vocabulary term base; test translations against agency-provided reference samples. Data migration and historical import: During this window, agencies can provide legacy Translation Memory (TMX files), historical bilingual documents, and existing terminology databases. Our engineering team imports these into the agency’s dedicated secure tenant so the AI model immediately leverages the agency’s historical linguistic preferences.

Day	Milestone
21–28	Integration configuration: connect to agency’s video conferencing platform (Zoom/Teams/Webex); Granicus integration if required
28–30	User acceptance testing: agency admins test full workflow; training sessions delivered; go-live sign-off
30+	Go-live; hypercare support (daily check-ins) for first 30 days

Technical Requirements (Agency Side): Minimal — no IT infrastructure changes required. Agencies have two deployment paths for live interpretation: - **Web app:** Any modern browser (Chrome, Firefox, Edge, Safari); microphone access; works immediately if meeting audio is playing through speakers - **Desktop app:** Downloadable for Mac/Windows; captures system audio directly; includes popout caption window for room display; no microphone needed

Attendees access interpretation via QR code or link — browser only, no app download required on their end. No on-premise hardware required.

Integration Support:

Interpret AI will develop and deliver the following integrations post-award. Our provider-agnostic connector architecture and REST API foundation make these integrations well within our engineering capability; timelines are committed deliverables, not estimates.

- Granicus: WebVTT caption injection and post-meeting transcript upload — delivered within 90 days of contract execution
- Zoom, Microsoft Teams, Webex, Google Meet: Live caption injection via each platform’s closed captioning API — delivered within 60 days of contract execution
- Emergency notification systems (Everbridge, CodeRED, Omnilert): Alert translation webhook for WS5 — delivered within 30 days of contract execution
- Civic Plus, Tyler Technologies, Accela: REST API integration and custom connectors — available on request; scoped and delivered per agency need

4.2 Training and Support

Initial Training: - Administrator training: 2-hour live virtual session covering platform setup, user management, glossary management, reporting dashboard - End-user training: 45-minute live virtual session covering how to join a meeting with interpretation and how to submit documents for translation - Training is recorded and made available in the agency’s knowledge base portal within 48 hours

Ongoing Training: - New feature training released with each quarterly update (video + release notes) - New staff onboarding: self-serve training modules available on-demand; live session available on request (no charge)

Support Tiers:

Tier	Description	Response SLA
Tier 1	General usage questions, account access, billing	4 business hours
Tier 2	Platform configuration, integration issues, translation quality disputes	8 business hours
Tier 3	System outages, data incidents, API failures	1 hour (24/7)
Emergency	WS5 emergency alert translation failures	15 minutes (24/7)

Support Channels: Email (support@bluejayai.net), web portal ticketing, dedicated Slack channel for Enterprise accounts, phone (24/7 for Tier 3/Emergency)

Dedicated Account Management: Each PPA is assigned a Customer Success Manager (CSM) who is their single point of contact for all non-technical questions, renewal discussions, and new service expansion. CSMs conduct quarterly business reviews.

4.3 Service Level Agreements (SLAs)

System Uptime:

Workstream	Uptime Guarantee	Measurement Period
WS1 — Live Voice Interpretation	99.95%	Monthly
WS2 — Document Translation Portal	99.9%	Monthly
WS4 — Platform	99.9%	Monthly
WS5 — Emergency Translation	99.99%	Monthly

Uptime excludes planned maintenance windows (maximum 4 hours/month; minimum 72 hours advance notice; scheduled between 2–6 AM local time for affected agencies).

Document Translation Turnaround (WS2):

Priority Tier	Turnaround	Description
Standard	3 business days	General documents up to 5,000 words
Priority	1 business day	Documents up to 3,000 words
Urgent	4 hours	Documents up to 1,000 words; surcharge applies
Emergency	2 hours	Critical public-facing documents; limited to 500 words; surcharge applies

Documents exceeding word limits at each tier are quoted separately. Turnaround clock starts upon receipt of complete source file.

Emergency Alert Translation (WS5): Sub-5-minute turnaround guaranteed for emergency alerts up to 500 words. 24/7/365. SLA measured from receipt of source content to delivery of translated output.

Service Credits for SLA Failures:

Failure	Credit
Monthly uptime below 99.5%	10% of that month's fees
Monthly uptime below 99.0%	20% of that month's fees
Monthly uptime below 98.0%	50% of that month's fees
Document delivery 25–50% beyond SLA	10% of job value
Document delivery >50% beyond SLA	25% of job value
Emergency alert turnaround exceeded	100% of alert fee

EVALUATION SECTION 5: PRICING

Workstreams Covered: 1, 2, 4, and 5

Workstream 1 — Live AI Voice Interpretation

Service	Unit	List Price	Cooperative Discount %
Live AI Interpretation — per meeting/hour (up to 50 participants)	Per hour	\$49.00	5%
Live AI Interpretation — per meeting/hour (51–200 participants)	Per hour	\$89.00	5%
Live AI Interpretation — per meeting/hour (201+ participants)	Per hour	\$149.00	5%
Monthly subscription — unlimited meetings (up to 10/month)	Per month	\$499.00	5%
Annual subscription — unlimited meetings	Per year	\$4,990.00	10%
Additional language channel (beyond base package)	Per language/meeting	\$15.00	5%
Post-Meeting Multilingual Transcription (AI-only)	Per Audio Minute	\$0.10	5%
Post-Meeting Multilingual Transcription (Human-Edited / 100% Accurate)	Per Audio Minute	\$2.50	5%

Workstream 2 — Document Translation

Service	Unit	List Price	Cooperative Discount %
AI document translation — standard (no human review)	Per word	\$0.02	5%
AI document translation — with human post-editing	Per word	\$0.15	5%
Emergency/rush translation (< 4 hour turnaround)	Per word	\$0.25	5%
Website content localization	Per page	\$5.00	5%
Annual document translation subscription	Per year	\$5,500.00	10%
Desktop Publishing (DTP) / Advanced Formatting (Adobe InDesign, Illustrator, etc.)	Per Hour	\$75.00	5%
Audio/Video Dubbing & Voiceover (AI Voice Generation)	Per Video Minute	\$10.00	5%
Video Subtitling / Closed Captioning (Human-Edited)	Per Video Minute	\$4.00	5%

Workstream 3 — OPI / VRI

Not Bidding. Interpret AI is specializing on Workstreams 1, 2, 4, and 5 to ensure highest-quality AI-native deployment.

Workstream 4 — AI Translation Platform / Software License

Service	Unit	List Price	Cooperative Discount %
SaaS Platform License — small agency (< 50,000 population)	Per year	\$5,000	10%
SaaS Platform License — mid-size agency (50k–300k population)	Per year	\$12,000	10%
SaaS Platform License — large agency (300k+ population)	Per year	\$25,000	10%

Service	Unit	List Price	Cooperative Discount %
API access	Per 1M characters	\$40.00	5%
On-premise/private cloud deployment (one-time setup + annual)	Lump sum	\$35,000	5%

Workstream 5 — Emergency & Public Safety Translation

Workstream 5 emergency translation services are billed at the standard Emergency/Rush per-word rates listed in the Workstream 2 table above (\$0.25/word, 5% cooperative discount), with sub-5-minute turnaround SLA for emergency alerts up to 500 words, 24/7/365.

Pricing Guarantee: All pricing is guaranteed for the initial 5-year contract term. Annual adjustments permitted only upon mutual written agreement and not to exceed CPI for that year, with 90 days advance written notice.

Most Favored Public Pricing: Interpret AI confirms that pricing in this proposal equals or beats pricing offered to any other cooperative program or individual entity with equal or lower anticipated volume.

APPENDIX SUBMISSIONS

Appendix	Status
Appendix A — Offer and Contract Signature Form	To be signed and uploaded to Indigo portal
Appendix B — Terms and Conditions Acceptance Form	To be signed and uploaded
Appendix C — Disclosures Federal/NJ	To be signed and uploaded
Appendix D — Sample Supplier Agreement	No redlines proposed; standard terms accepted
Appendix E — Language Coverage Matrix	Separate XLSX attached
Appendix F — Security & Data Privacy Certification	Completed and attached
Appendix G — AFI Master Agreement	No redlines proposed; standard terms accepted
Appendix H — Contract Sales Report Template	Acknowledged; will report monthly by 10th
Appendix I — Insurance Certificate	CGL: Hiscox P105.122.567.2, \$2M/\$2M; Workers Comp: State Fund 9383216-2025, \$1M; Tech E&O / Cyber Liability: Policy will be bound and certificate provided prior to contract execution, \$2M/\$5M
Appendix J — Platform Integration Reference Sheet	Provided below
Appendix K — SLA Template	Completed and attached
Appendix L — AI Accuracy Disclosure	Signed and attached

Appendix J — Platform Integration Reference:

Platform	Integration Type	Status
----------	------------------	--------

Platform	Integration Type	Status
Zoom	Live caption injection (CC API), post-meeting transcript	Within 60 days of contract execution
Microsoft Teams	Live caption injection, transcript	Within 60 days of contract execution
Webex	Live caption injection	Within 60 days of contract execution
Google Meet	Live caption injection	Within 60 days of contract execution
Granicus	WebVTT caption file injection; post-meeting transcript upload	Within 90 days of contract execution
Everbridge	Emergency alert translation webhook	Within 30 days of contract execution
CodeRED / Omnilert / Rave	Emergency alert REST API	Within 30 days of contract execution
REST API (any platform)	Full translation/interpretation API with OAuth 2.0	Within 30 days of contract execution
SSO	SAML 2.0, Azure AD, Okta	Within 30 days of contract execution

Proposal prepared May 11, 2026 — Bluejay AI Technologies Inc Point of Contact: Joseph Barrientos Liba — support@bluejayai.net — (978) 799-7069