




II ElevenLabs

Service Definition Document

AI Voice Generation & Text-to-Speech

G-Cloud 15 – Lot 2b: Cloud Software

1. Service Overview

2. Core Speech Synthesis Architecture

3. Key Capabilities

3.1 Voice Selection and Voice Creation

3.2 Multilingual Speech Output

3.3 Streaming and Low-Latency Output

3.4 Output Formats for Delivery Requirements

3.5 Repeatability and Consistency Controls

3.6 Long-Form and High-Volume Generation Patterns

4. Onboarding, Support and Service Levels

4.1 Onboarding

4.2 Support and After-Sales Service

4.3 Service Levels and Performance

- 4.4 Service Constraints
- 5. Hosting, Deployment and Data Residency
 - 5.1 Hosting Model
 - 5.2 Data Residency
- 6. Security & Privacy
 - 6.1 Security Governance
 - 6.2 Information Security Controls
 - 6.3 Vulnerability Management
 - 6.4 Protective Monitoring
 - 6.5 Incident Management
 - 6.6 Access Control
 - 6.7 Certifications and Assurance
- 7. Backup, Restore and Disaster Recovery
 - 7.1 Backup and Restore
 - 7.2 Disaster Recovery and Business Continuity
- 8. Operations, Maintenance and Updates
 - 8.1 Operational Monitoring and Observability
 - 8.2 Product-Specific Monitoring and Quality Evaluation
 - 8.3 Logging, Data Collection and Audit Support
 - 8.4 Service Updates and API Versioning
 - 8.5 Outage and Maintenance Management
 - 8.6 Continuous Improvement
- 9. Service Constraints
 - 9.1 Usage Constraints
 - 9.2 Customisation
- 10. Exit Management and Data Access
 - 10.1 Data Ownership
 - 10.2 Service Offboarding and Data Export
 - 10.3 Termination, Deletion and Wind-Down
- 11. Technical Requirements

1. Service Overview

ElevenLabs AI Voice Generation & Text-to-Speech (TTS) is a managed, cloud-based speech synthesis service that converts text into spoken audio. It is designed for public-sector and enterprise use cases such as accessibility services, public information and announcements, e-learning content, and voice-enabled digital applications.

The service is delivered as software-as-a-service (SaaS) through web-based tools and APIs. Customers generate audio without managing model hosting, scaling, or audio processing infrastructure. The service supports both non-real-time generation (producing complete audio files) and streaming generation (progressively producing audio for lower perceived latency in interactive experiences).

This service provides speech synthesis only. Speech recognition, conversational reasoning, dialogue management, and call handling are not included.

2. Core Speech Synthesis Architecture

The AI Voice Generation & TTS service coordinates the following components to deliver speech output reliably and at scale:

- **Text Input and Request Controls:** Customers submit text and synthesis settings via API or the web interface. Requests may include configuration parameters to control voice selection, output stability/consistency, and repeatability.
- **Voice Profile Selection and Configuration:** Customers can use pre-built voices or create authorised voice profiles using supported voice creation methods. Voice configuration is applied independently of downstream delivery channels.
- **Speech Synthesis and Audio Rendering:** The service generates speech audio from text using ElevenLabs speech synthesis models. Audio can be returned as a complete file or streamed in chunks for progressive playback.
- **Output Encoding and Format Handling:** The service supports multiple output encodings and sampling rates to match delivery environments, including web and mobile media formats and formats commonly used in voice delivery systems (for example, 8kHz telephony-compatible codecs), without providing telephony connectivity or call orchestration.

3. Key Capabilities

3.1 Voice Selection and Voice Creation

Customers can select from a catalogue of voices or create voice profiles under supported mechanisms. Voice creation options include rapid voice creation from short samples and higher-assurance methods intended for authorised voice replicas. Voice governance requirements (such as consent and verification for certain voice creation methods) apply where relevant to the selected capability.

3.2 Multilingual Speech Output

The service supports multilingual speech synthesis across a broad set of languages. Customers can deploy consistent voice outputs across different languages and service lines by selecting appropriate voices and synthesis configurations for each target language and locale.

3.3 Streaming and Low-Latency Output

The service supports streaming output for use cases where audio should begin playing before the full synthesis completes. This reduces perceived latency for end users and supports interactive applications that require responsive speech output.

3.4 Output Formats for Delivery Requirements

Audio can be generated in multiple formats suitable for different operational contexts, including compressed media formats and raw PCM formats. Customers can configure output parameters such as encoding and sample rate to match their delivery environment requirements.

3.5 Repeatability and Consistency Controls

Speech synthesis is probabilistic and may vary slightly between requests. Where consistent outputs are required, customers can use supported configuration parameters (such as request seeding and stability controls) to increase repeatability and standardise delivery across runs,

while recognising that minor variation can still occur.

3.6 Long-Form and High-Volume Generation Patterns

For longer content, the service supports generation patterns that maintain quality and continuity, including chunking and streaming strategies and optional contextual parameters to improve prosody across segments. This supports production use cases such as long-form narration and large-scale content generation.

4. Onboarding, Support and Service Levels

4.1 Onboarding

The ElevenLabs Text-to-Speech API is delivered as a managed, cloud-based SaaS offering.

Customers receive access to standard onboarding resources, including:

- API documentation and technical references
- Integration guides and examples
- Online documentation describing supported features and usage

4.2 Support and After-Sales Service

Support and maintenance are provided for the duration of the Call-Off Contract.

Standard Support (Included)

Standard support is included as part of the core service offering and provides:

- Access to technical support for service-related issues and usage queries
- Email- or ticket-based support channels
- Support availability during published support hours, excluding public holidays
- Incident investigation, fault resolution, and guidance on the use of existing service functionality

Optional Enhanced Support Services

In addition to standard support, ElevenLabs offers enhanced customer success and support services, which may be procured as an optional add-on. These services are not required to use the core service and do not alter the service availability commitments or service credit remedies applicable under the Service Level Agreement. Where purchased, enhanced support services may include defined support response time targets and additional engagement channels, as described in the Pricing Document.

4.3 Service Levels and Performance

The Text-to-Speech API is subject to availability commitments set out in the applicable Service Level Agreement.

For the TTS API, ElevenLabs commits to **99.9% monthly availability**, measured on a per-minute basis. Availability is measured on a 24x7 basis and excludes scheduled maintenance and standard SLA exclusions.

Performance characteristics may vary depending on request parameters, configuration, and network conditions.

Full definitions, measurement methodology, exclusions, and remedies are set out in the Call-Off Contract and SLA.

4.4 Service Constraints

ElevenLabs AI Voice Generation & TTS is delivered as a standardised, multi-tenant SaaS service. Customers cannot obtain bespoke development or work-for-hire under this service. The service requires internet connectivity and supported clients for web or API access. Speech synthesis performance and latency may vary with configuration, model choice, input length, output format, and network conditions. Planned maintenance may occur to maintain security and stability, with disruption minimised where possible. ElevenLabs may limit or suspend access to protect service integrity or address security risks. Customers must not submit PHI without a valid agreement, or special-category data unless permitted by applicable law.

5. Hosting, Deployment and Data Residency

5.1 Hosting Model

The service operates on a one-to-many multi-tenant SaaS model with data segregated at the application layer.

- **Cloud Provider:** The service is hosted on Google Cloud Platform (GCP) in secure data centres.
- **Infrastructure:** ElevenLabs utilises technology service suppliers to provide hosting solutions and support technology, maintaining responsibility for supplier performance.

5.2 Data Residency

ElevenLabs provides options regarding the processing location of data:

- **Standard Residency:** By default, data is processed in US-based regions.
- **EU Residency:** EU-based environments are available for buyers requiring specific residency controls. For customers requiring EU data residency, service user data is stored and processed within the European Economic Area (EEA).
- **Subprocessors:** A list of subprocessors is available via the ElevenLabs compliance portal.

6. Security & Privacy

6.1 Security Governance

ElevenLabs operates a formal information security and privacy management framework designed to protect the confidentiality, integrity, and availability of Customer Content. Security governance includes defined ownership, reporting structures, and regular review of policies and controls, which are documented, maintained and made publicly available at <https://compliance.elevenlabs.io>. All personnel are required to formally acknowledge relevant security policies as part of onboarding and on an ongoing basis.

6.2 Information Security Controls

The service is protected by technical and organisational measures covering:

- Access control and authentication

- Data protection and encryption
- Logging, monitoring, and audit
- Secure software development practices
- Change and configuration management
- Vendor and asset management
- Privacy and data protection compliance
- AI governance and risk management

Security controls are reviewed and updated to address changes in risk, regulation, and technology.

6.3 Vulnerability Management

ElevenLabs operates a vulnerability management programme to identify, assess, and remediate security risks. This includes:

- Continuous assessment of potential threats using internal and external intelligence sources
- Regular vulnerability scanning
- Independent penetration testing conducted at least annually
- Risk-based prioritisation and remediation of findings

6.4 Protective Monitoring

The service is continuously monitored to detect security events, misuse, and operational anomalies. Monitoring includes:

- Detection of unusual access patterns or behaviour
- Automated and manual alerting mechanisms
- Investigation of potential security incidents
- Abuse and misuse detection

ElevenLabs may suspend or limit access to the service where necessary to protect platform integrity or address a security risk.

6.5 Incident Management

ElevenLabs maintains predefined incident response processes for security and operational events. Customers may report incidents through established support channels. Where applicable, incident notifications and reports are provided in accordance with contractual and regulatory requirements.

6.6 Access Control

Access to management interfaces and support systems is restricted using role-based access controls and authentication mechanisms. Access is limited to authorised personnel and subject to logging and monitoring. Customer data is accessed only where required to provide support.

6.7 Certifications and Assurance

The service is supported by recognised security and compliance certifications and attestations, including:

- ISO/IEC 27001
- SOC 2 Type II
- Cyber Essentials and Cyber Essentials Plus
- PCI DSS
- CSA STAR (Level 1)
- Additional privacy and AI governance standards

These certifications are maintained through ongoing audit and assurance activities.

7. Backup, Restore and Disaster Recovery

7.1 Backup and Restore

- **Automated Backups:** ElevenLabs performs regular automated backups of platform configuration and retained data as required for the operation of the SaaS service.
- **Data Integrity:** Measures are in place to protect the security and integrity of Customer Content collected, accessed, stored, or received.

7.2 Disaster Recovery and Business Continuity

The platform is designed for high availability using cloud-native redundancy:

- **Resilience:** The solution is hosted in GCP data centres, leveraging availability zones to maintain service continuity.
- **Force Majeure:** The service includes provisions for events beyond reasonable control (e.g., internet failures, public health crises), during which performance obligations may be delayed.

8. Operations, Maintenance and Updates

8.1 Operational Monitoring and Observability

ElevenLabs operates continuous monitoring and observability for the Service to support reliable, secure, and compliant operation in production environments.

Operational monitoring may include:

- Service availability and performance indicators
- Request and processing metrics
- Error rates and failure conditions
- Abuse, misuse, and anomalous usage patterns

Monitoring is designed to support service assurance, incident response, and operational governance without exposing underlying infrastructure or proprietary implementation details.

8.2 Product-Specific Monitoring and Quality Evaluation

For the Text-to-Speech service, monitoring and evaluation focus on speech synthesis quality and stability.

This may include:

- Synthesis latency and availability
- Output consistency and voice stability
- Detection of synthesis failures or degraded audio output
- Usage patterns across voices and languages

Monitoring supports operational reliability and continuous improvement of speech quality without requiring customer-side infrastructure management.

8.3 Logging, Data Collection and Audit Support

The Service provides mechanisms to collect and retain operational logs and results in support of audit, investigation, and service improvement activities.

Logged data may include:

- Request metadata and timestamps
- Processing outcomes and system-generated events
- Input and output artefacts, where applicable and enabled by the Customer

Customers may access available logs and results via APIs or web-based tools, subject to configured retention policies and applicable data protection requirements.

8.4 Service Updates and API Versioning

ElevenLabs modifies the Services, including features and functionality, on an ongoing basis as part of the SaaS delivery model.

- Updates are applied centrally and made available without local installation.
- Updates will not result in a material degradation of the overall functionality or performance of the Services used by the Customer.
- APIs are versioned where applicable.
- At least 90 days' prior notice is provided before deprecating older API versions.

8.5 Outage and Maintenance Management

Planned maintenance is periodically required to ensure the security, stability, and performance of the Service.

- ElevenLabs endeavours to minimise disruption and schedule maintenance outside peak usage where practical.
- Service availability and incident information are published via the public service status page.

- Incidents are classified and managed based on severity and impact.
- In exceptional circumstances, including high-risk security events, access to specific accounts or features may be temporarily restricted to protect platform integrity.

8.6 Continuous Improvement

Operational insights derived from monitoring, evaluation, and customer feedback are used to inform continuous improvement of the Service. Enhancements to reliability, performance, and quality are delivered through the standard update mechanisms described above.

9. Service Constraints

9.1 Usage Constraints

- **Prohibited Use:** Customers must comply with the ElevenLabs Prohibited Use Policy.
- **Sensitive Data:** Customers must not provide Input that includes Protected Health Information (PHI) unless an agreement is executed, or special category data (e.g., biometric for identification, political opinions) unless permitted by applicable law.
- **Reverse Engineering:** Customers may not reverse engineer, decompile, or attempt to discover source code or underlying components of the Services or Models.

9.2 Customisation

The service is a standardised SaaS offering. Custom development or "work-for-hire" is not available. Customisation is achieved through the configuration of Agents, Voices, and Workflows within the standard tooling.

10. Exit Management and Data Access

10.1 Data Ownership

- **Input Rights:** The Customer retains all rights in Input (data transmitted to the system).
- **Output Rights:** As between the Customer and ElevenLabs, the Customer owns all rights in Output (content generated by the Services). ElevenLabs does not train its models on Enterprise Customer Output, except as may be required to provide certain requested services to the Customer.

10.2 Service Offboarding and Data Export

Upon termination or expiry of the Call-Off Contract, or during the contract term, customers may retrieve their Input and Output data on a self-serve basis.

- **Export Mechanisms:** Data retrieval is supported via standard export mechanisms provided by the platform, including APIs and web-based tools.
- **Formats:** Data is typically available in standard formats (e.g., JSON for configuration/logs, Audio files for output).
- **Documentation:** ElevenLabs provides documentation describing available export formats and data access mechanisms. No bespoke migration, data transformation, or third-party integration services are included as part of offboarding.

10.3 Termination, Deletion and Wind-Down

- **Cessation of Use:** Upon termination, the Customer must immediately cease all use of the Services.
- **Data Deletion:** Customer Content is deleted from the service within thirty (30) days of contract termination or expiry, unless the Customer explicitly agrees to retention beyond this period. Data deletion processes are designed to ensure secure and irreversible removal.
- **Wind-Down:** Specific wind-down periods may apply for partners/resellers distributing the service to End Users, ensuring a managed exit up to one year from expiration.
- **Sanitisation:** Industry-standard data sanitisation processes are applied to storage media, including secure erasure or physical destruction where applicable.

11. Technical Requirements

To access and use the ElevenLabs AI Voice Generation & Text-to-Speech service, the Customer is responsible for obtaining and maintaining the following:

- **Equipment:** Hardware, servers, and operating systems necessary to connect to the cloud service.
- **Network:** Internet connectivity and networking capability.
- **Web Access:** A modern web browser for the administration dashboard.
- **API Integration:** Capabilities to integrate via REST API for programmatic access (if required).