AIR GAP VOICE

SCIF COMPLIANCE GUIDE

April 2026

PUBLIC RELEASE

This guide evaluates Airgap Voice, a fully offline, on-device speech-to-text application, against the applicable security controls, directives, and operational requirements governing classified information processing environments. It addresses NIST SP 800-53 control mapping, the RMF accreditation pathway, air-gapped deployment procedures, TEMPEST considerations, and residual risk assessment for deployment at CUI through TS/SCI classification levels.

Author: Marvin F. L. Hansen
Contact: marvin.hansen@airgapvoice.com
Web: www.airgapvoice.com

Table of Contents

1. Purpose and Scope

1.1 Problem Statement

Automated transcription has become an operational necessity across the Department of Defense (DoD) and the Intelligence Community (IC). Personnel routinely require speech-to-text capability for meeting documentation, debriefing capture, after-action reporting, and intelligence product generation. However, the predominant transcription solutions on the market require network connectivity and transmit audio data to vendor-operated servers for processing.

This architecture is fundamentally incompatible with classified environments. Any system operating within a Sensitive Compartmented Information Facility (SCIF) or processing information at the Controlled Unclassified Information (CUI), Secret, or Top Secret/Sensitive Compartmented Information (TS/SCI) level must comply with stringent data handling, boundary protection, and emanation security requirements that categorically prohibit the transmission of processed data to external servers.

The purpose of this guide is to evaluate Airgap Voice against the applicable security controls, directives, and operational requirements governing classified information processing environments, and to assess its viability as a transcription tool within such environments.

1.2 Scope of Applicability

This assessment addresses deployment scenarios across the following classification levels:

- Controlled Unclassified Information (CUI) per 32 CFR Part 2002
- Secret / Collateral classified environments
- Top Secret / Sensitive Compartmented Information (TS/SCI) environments
- Special Access Program (SAP) facilities where applicable

The assessment is limited to the software layer. Physical security, personnel security, and facility accreditation are outside the scope of this document, though relevant intersections (e.g., TEMPEST, physical access controls) are addressed where they bear on the software deployment.

1.3 Reference Framework

This guide evaluates compliance against the following primary frameworks:

- NIST Special Publication 800-53 Revision 5: Security and Privacy Controls for Information Systems and Organizations
- CNSSI 1253: Security Categorization and Control Selection for National Security Systems
- ICD 503: Intelligence Community Information Technology Systems Security Risk Management
- DISA Security Technical Implementation Guides (STIGs): macOS platform-specific guidance
- NIST SP 800-37 Revision 2: Risk Management Framework for Information Systems and Organizations

2. Applicable Standards and Directives

2.1 Governing Standards

The following table identifies the primary standards, directives, and guidelines applicable to the evaluation and accreditation of transcription software in classified environments.

Standard / Directive	Applicability	Relevance to Transcription Software
NIST SP 800-53 Rev 5	All federal information systems	Defines baseline security controls; SC, AC, AU, and SI families directly applicable
CNSSI 1253	National Security Systems (NSS)	Establishes control baselines by confidentiality impact level for NSS
ICD 503	IC information systems	Governs risk management and authorization for IC systems; defines ATO requirements
DCID 6/3 (legacy)	IC systems (superseded by ICD 503)	Legacy reference; some accreditation artifacts still reference DCID 6/3 language
DISA macOS STIG	DoD macOS endpoints	Platform-specific hardening requirements for macOS in DoD environments
NIST SP 800-37 Rev 2	All federal systems undergoing authorization	Defines the RMF process; governs ATO lifecycle
CNSS Policy 15 (CNSSP-15)	National Security Systems	Governs use of public standards for securing NSS
DoD Instruction 8510.01	DoD information systems	Implements RMF for DoD; defines roles and authorization process

2.2 TEMPEST and EMSEC Requirements

Systems processing classified information within SCIFs are subject to TEMPEST requirements as defined in CNSSAM TEMPEST/01-13 and related classified annexes. Key considerations include:

- Electromagnetic emanation profiles of processing hardware
- Acoustic emanation from hardware components (fans, speakers, drives)
- TEMPEST zone classification and countermeasure requirements
- RED/BLACK separation requirements for classified data processing

2.3 Cryptographic Requirements

FIPS 140-2 and FIPS 140-3 establish cryptographic module validation requirements for systems protecting sensitive and classified information. Airgap Voice does not implement, require, or invoke cryptographic functions as part of its transcription processing. Audio data is processed entirely in volatile memory without encryption because encryption is unnecessary when data never leaves the device boundary and is never written to persistent storage.

If the deployment environment requires encrypted storage for the application binary itself (e.g., FileVault on macOS), the host operating system's FIPS-validated cryptographic module (Apple corecrypto, FIPS 140-3 certificate #4701) provides that function at the platform level, external to the application.

2.4 FedRAMP Applicability

FedRAMP (Federal Risk and Authorization Management Program) governs the security assessment of cloud services used by federal agencies. FedRAMP authorization is predicated on a cloud service model: the provider operates infrastructure that processes federal data.

Airgap Voice does not operate as a cloud service. It performs no network communication, maintains no cloud infrastructure, and transmits no data to any external system. There is no service provider boundary. Consequently, FedRAMP authorization is not applicable to this software. The relevant accreditation pathway is the Risk Management Framework (RMF) as applied to the host information system on which the software is installed.

3. Threat Analysis for Transcription Systems

Before assessing any specific product, it is necessary to enumerate the threat vectors inherent to transcription systems when deployed in classified environments. This analysis informs the security control evaluation in subsequent sections.

3.1 Data Exfiltration via Network

The most significant threat posed by commercial transcription solutions is the transmission of audio data — potentially containing classified information — to vendor-operated cloud servers. This threat is not theoretical; it is the standard operating model for major cloud dictation APIs. Audio data sent to these services crosses organisational boundaries, jurisdictional boundaries, and in some cases national boundaries.

For classified environments, this threat vector is an absolute disqualifier. NIST SP 800-53 control SC-7 (Boundary Protection) requires monitoring and control of communications at the system boundary. Any software that initiates outbound network connections carrying processed data cannot satisfy SC-7 in a classified enclave.

3.2 Data at Rest Exposure

Transcription systems that write audio recordings, intermediate processing files, or transcription output to disk create persistent classified data artifacts that must be managed throughout their lifecycle. Temporary files, cache directories, swap space, and log files can all contain fragments of classified content. This increases the attack surface and complicates data sanitisation procedures per NIST SP 800-88.

An ideal transcription system for classified environments would process audio exclusively in volatile memory, producing only the final text output that the user explicitly directs to a destination (e.g., insertion at the cursor position in a classified document).

3.3 Supply Chain Compromise

Many AI transcription tools require runtime model downloads from vendor servers, periodic updates, or license activation against external servers. Each of these mechanisms introduces a supply chain dependency and a potential vector for compromise. A model file downloaded at runtime could be tampered with in transit or at the source. An update mechanism provides a persistent channel through which malicious code could be introduced.

For deployment in air-gapped environments, the software must be fully self-contained: all models bundled within the application binary, no runtime downloads, no update mechanism that requires network connectivity, and no license activation that contacts a vendor server.

3.4 Covert Channel Analysis

Covert channels — both storage and timing — must be assessed for any software operating in a classified environment (per NIST SP 800-53 SC-31). For a transcription application, the relevant covert channel considerations include:

- Acoustic emanation: Does the software cause the hardware to produce sounds that encode information?
- Electromagnetic emanation: Does processing activity create detectable EM signatures correlated with input data?
- Network covert channels: Can the software establish unauthorised communication via DNS, ICMP, or other protocols?
- Filesystem covert channels: Does the software write data to unexpected locations accessible by other processes?

For a software application that performs no network I/O, does not write to disk, and operates within an application sandbox, the covert channel risk is substantially reduced but not eliminated. Electromagnetic and acoustic emanation remain considerations addressed in Section 8.

3.5 Insider Threat Considerations

No software solution eliminates insider threat risk. A cleared individual with physical access to the endpoint could, in principle, copy transcribed text, reconfigure the system, or abuse the microphone access. However, the transcription software itself should not amplify insider threat risk beyond the baseline established by the endpoint's operating system and access controls.

Airgap Voice's architecture (no persistent storage of audio, no network capability, no export mechanism) means it does not create new exfiltration paths that an insider could exploit. The transcribed text exists only at the cursor insertion point within whatever application the user is working in, subject to that application's own security controls.

4. Architecture Assessment: Airgap Voice

4.1 Processing Model

Airgap Voice performs all speech recognition processing locally on the host device. The application is built exclusively for Apple Silicon (ARM64 architecture) and leverages the Metal GPU framework and Neural Engine co-processor for inference acceleration. No x86 emulation layer is involved; the application runs natively on ARM64.

The processing pipeline is entirely self-contained. Audio captured from the device microphone is passed through locally-stored inference models that perform speech-to-text conversion. At no point in this pipeline does data leave the device. There is no network dependency of any kind: no API calls, no cloud inference, no telemetry, no analytics, no crash reporting, no update checks.

> **Verification: The zero-network claim is empirically verifiable using standard system monitoring tools available on macOS: tcpdump, lsof -i, nettop, and Activity Monitor → Network tab. These procedures can be incorporated into the Security Assessment Plan as part of the RMF authorization process.**

4.2 Data Flow Analysis

Stage	Data Location	Persistence	Network Activity
1. Audio Capture	Microphone hardware buffer	None — real-time stream	None
2. Audio Processing	Volatile memory (RAM)	None — in-memory only	None
3. Model Inference	Volatile memory (GPU/Neural Engine)	None — in-memory only	None
4. Text Output	Cursor insertion in active application	User-controlled	None
5. Buffer Cleanup	Memory zeroed via memset_s	Securely erased	None

At Stage 5, audio buffers are securely zeroed using memset_s, which the compiler is prohibited from optimising away per ISO/IEC 9899:2011 (C11) Annex K. This provides assurance that audio data does not persist in memory beyond its immediate processing window.

4.3 Model Delivery

The speech recognition models used by Airgap Voice are bundled within the application binary itself (the macOS .app bundle). There is no runtime model download, no model update mechanism that contacts external servers, and no dynamic model loading from user-accessible directories.

This is a critical architectural property for classified deployment. It eliminates the supply chain risk associated with runtime model acquisition and ensures that the application's behaviour is fully deterministic and auditable from the point of initial transfer to the classified network.

4.4 Licensing Model

> **This section addresses a critical operational requirement for air-gapped deployment.**

Airgap Voice uses a site license model specifically designed for environments where network connectivity is unavailable or prohibited:

- Site license: A single procurement action authorises deployment across all machines within the licensed organisation or facility.
- No license server: The software does not contact, require, or attempt to reach any license validation server.
- No license keys: The software does not require entry of a license key, activation code, or serial number during installation or at any point during operation.
- No activation: There is no activation step, online or offline. The software is fully functional immediately upon installation.
- No phone-home: The software makes zero network connections for any purpose, including license validation. This is verifiable using standard network monitoring tools.
- No expiration check: The license does not enforce time-based expiration through network or local clock validation.

This licensing model eliminates a class of deployment obstacles that frequently impede the introduction of commercial software into air-gapped classified networks. The deployment process is: transfer the application binary via approved secure media, install, and operate. No additional infrastructure, no license management servers, and no network exceptions are required.

4.5 Permission Model

Airgap Voice requests exactly two system permissions:

- Microphone Access: Required for audio capture. Granted via macOS TCC (Transparency, Consent, and Control) framework. Revocable at any time via System Settings.
- Accessibility Access: Required for text insertion at the cursor position in the active application. Granted via macOS TCC framework. Revocable at any time via System Settings.

The application does not request network access, file system access beyond its sandbox container, contacts, calendar, location, camera, or any other system permission. The minimal permission footprint reduces the attack surface and simplifies the security assessment.

4.6 Build Hardening

Airgap Voice is built with the following macOS platform security features enabled:

- App Sandbox: Enforced by macOS. Restricts file system access, network access, and inter-process communication to explicitly declared entitlements.
- Hardened Runtime: Enforced by macOS. Prevents code injection, dylib hijacking, and debugging by unauthorised processes.
- Code Signing: Application is signed with a valid Apple Developer ID certificate. Signature integrity can be verified using codesign --verify --deep --strict.
- Notarisation: Application has been submitted to and approved by Apple's notarisation service, which performs automated malware scanning. Verifiable via spctl --assess --type exec.
- ARM64-only: The binary contains only ARM64 (Apple Silicon) code. No x86_64 slice is present, eliminating the Rosetta 2 translation layer and its associated attack surface.

5. NIST SP 800-53 Control Mapping

The following table maps selected NIST SP 800-53 Rev 5 security controls to the Airgap Voice architecture. This mapping is intended to support the Security Assessment Plan (SAP) and inform the Security Assessment Report (SAR) during the RMF authorisation process.

Control ID	Control Name	Airgap Voice Implementation	Assessment Method
AC-3	Access Enforcement	macOS TCC framework enforces per-permission access. Microphone and Accessibility require explicit user/admin grant. App Sandbox enforces file system and IPC restrictions.	Verify TCC grants via tccutil; review entitlements via codesign --display --entitlements
AC-4	Information Flow Enforcement	Data flows exclusively from microphone to volatile memory to cursor insertion. No outbound network flow. No persistent storage flow. Enforced by App Sandbox network entitlement absence.	tcpdump capture during operation; lsof -i verification; entitlements review
AU-2	Event Logging	Application does not generate its own audit logs to avoid creating persistent classified data artifacts. macOS Unified Logging captures process lifecycle events at the OS level.	Review Unified Log for process events; verify absence of application-level log files in sandbox container
SC-7	Boundary Protection	The application initiates zero network connections. App Sandbox does not include the network.client entitlement. The network boundary is not crossed.	tcpdump full capture during operation; codesign entitlements review confirming absence of network entitlement
SC-8	Transmission Confidentiality and Integrity	Not applicable. No data is transmitted. The application performs no network I/O.	Verify via tcpdump, lsof, nettop
SC-28	Protection of Information at Rest	Audio data is never written to disk. Processing occurs exclusively in volatile memory. Buffers are zeroed using memset_s after use. No cache files, no temp files, no log files containing audio data.	File system monitoring during operation (fs_usage); review sandbox container for absence of audio data artifacts
SI-2	Flaw Remediation	Application updates are delivered via secure media transfer for air-gapped deployments. Code signature and notarisation provide integrity verification for each update.	Verify code signature on update binary; compare hash against manifest provided via out-of-band channel
SI-3	Malicious Code Protection	Apple notarisation includes automated malware scanning. Hardened Runtime prevents code injection. App Sandbox prevents unauthorised IPC.	spctl --assess --type exec; codesign --verify --deep --strict

Control ID	Control Name	Airgap Voice Implementation	Assessment Method
SI-4	System Monitoring	Standard macOS monitoring tools (Activity Monitor, tcpdump, lsof, fs_usage) provide full observability into application behaviour. No obfuscation, anti-debugging, or monitoring evasion techniques are employed.	Continuous monitoring via macOS system tools; periodic spot checks per CM plan
SC-31	Covert Channel Analysis	No network covert channels (no network access). No filesystem covert channels (sandbox isolation). Acoustic/EM emanation addressed in Section 8.	Network monitoring; filesystem monitoring; TEMPEST assessment per Section 8
SC-4	Information in Shared System Resources	Audio buffers are zeroed via memset_s after processing. This function cannot be optimised away by the compiler, providing assurance of secure buffer clearing.	Code review; dynamic analysis confirming memset_s invocation

6. Accreditation Pathway

6.1 Risk Management Framework (RMF) Integration

Airgap Voice should be assessed as a software component of the host information system on which it is deployed, not as a standalone system requiring its own authorisation boundary. The application does not establish an independent network presence, does not maintain persistent data stores, and does not operate services.

Step 1: Categorize

The information system's existing categorisation (per FIPS 199 and CNSSI 1253) applies to Airgap Voice as a component. No separate categorisation is required. The software processes audio data at the same classification level as the environment in which it operates.

Step 2: Select Controls

The control baseline established for the host system applies. The control mapping in Section 5 identifies how Airgap Voice satisfies or is not applicable to relevant controls. No additional controls are introduced beyond those already required by the host system's baseline.

Step 3: Implement Controls

Many controls are inherently satisfied by the application's architecture (e.g., SC-7 is satisfied by the absence of network capability). Others are satisfied by the host operating system (e.g., AC-3 via macOS TCC). The implementation evidence is the architecture itself, verifiable through the assessment methods identified in Section 5.

Step 4: Assess Controls

The Security Control Assessor (SCA) should execute the assessment methods identified in the Section 5 control mapping table. Suggested test procedures are detailed in Section 6.3.

Step 5: Authorize

The Authorising Official (AO) reviews the Security Assessment Report (SAR) and Plan of Action and Milestones (POA&M) to issue an Authorisation to Operate (ATO) for the host system inclusive of Airgap Voice as a software component.

Step 6: Monitor

Continuous monitoring requirements for Airgap Voice are minimal given the application's static architecture. Periodic verification of code signature integrity, entitlement review, and network activity spot checks are sufficient.

6.2 Authorization to Operate (ATO) Considerations

Airgap Voice is well-suited for inclusion in an existing ATO or a new ATO package for the following reasons:

- It introduces no new network connections or boundary crossings.
- It introduces no new persistent data stores requiring separate data handling procedures.
- It requires no additional infrastructure (no servers, no databases, no license managers).
- Its behaviour is fully observable and verifiable using standard macOS system tools.
- Its code signature and notarisation provide a verifiable chain of integrity from the developer.

6.3 Suggested Test Procedures for Security Assessment

Test 1: Network Activity Verification

- Start tcpdump on all interfaces prior to launching application
- Operate application through full transcription workflow (start, dictate, stop, review output)
- Terminate tcpdump and review capture
- Expected result: Zero packets originating from or destined to the application process

Test 2: File System Activity Verification

- Start fs_usage filtered to the application process prior to launch
- Operate application through full transcription workflow
- Review fs_usage output for file write operations
- Expected result: No write operations to files containing audio data or transcription content outside the application sandbox container

Test 3: Code Signature and Entitlements Verification

```
codesign --verify --deep --strict /Applications/Airgap\ Voice.app
codesign --display --entitlements :- /Applications/Airgap\ Voice.app
spctl --assess --type exec /Applications/Airgap\ Voice.app
```

Expected result: Valid signature, notarisation confirmed, entitlements limited to microphone and accessibility.

Test 4: Open Socket Verification

```
lsof -i -p $(pgrep "Airgap Voice")
```

Expected result: No open network sockets.

6.4 Evidence Collection

The following artefacts should be collected and retained as part of the authorisation package:

- tcpdump capture file (.pcap) from Test 1
- fs_usage log from Test 2
- codesign and spctl output from Test 3
- lsof output from Test 4
- Screenshot of application entitlements
- Hash (SHA-256) of the application binary as installed on the classified system
- Copy of the site license agreement confirming no activation server requirement

7. Deployment Procedures for Air-Gapped Networks

7.1 Secure Media Transfer

Airgap Voice is delivered as a standard macOS application bundle (.app). For deployment to air-gapped networks, the application must be transferred via approved secure media in accordance with the facility's data transfer procedures. Typical approved media include:

- Write-once optical media (CD-R/DVD-R) with hash verification
- Approved USB storage devices per facility SOP
- Secure file transfer via cross-domain solution (CDS) where available and approved

The SHA-256 hash of the application bundle should be computed on the source system and verified on the destination system after transfer. The hash should be communicated via an out-of-band channel (e.g., printed manifest, separate classified communication) to provide integrity assurance independent of the transfer media.

7.2 Installation Without Network Connectivity

Installation of Airgap Voice on an air-gapped system requires no network connectivity at any stage. The installation procedure is:

- Copy the Airgap Voice.app bundle from approved media to /Applications/
- Verify code signature: codesign --verify --deep --strict /Applications/Airgap\ Voice.app
- Grant microphone and accessibility permissions via System Settings > Privacy & Security (requires administrator credentials)
- Launch application and verify operation with a test transcription

There is no installer that contacts external servers. There is no activation step. There is no license key to enter. There is no registration process. The software is fully operational immediately upon placement in the Applications directory and permission grant.

7.3 No License Activation Required

> **The site license model eliminates all network-dependent licensing mechanisms. Once procured, the software may be deployed to any number of machines within the licensed organisation without per-machine activation, license key entry, or server validation.**

No license management infrastructure is required. No network exceptions or firewall rules are needed. No license server needs to be accredited, authorised, or maintained within the authorisation boundary.

7.4 Configuration Management

Upon installation, the following baseline documentation should be established and maintained:

- Application version number and build identifier
- SHA-256 hash of the application bundle
- Code signing certificate details (issuer, subject, expiry)
- Installed entitlements (via codesign --display --entitlements)
- TCC permissions granted (microphone, accessibility)
- Host system identifier (serial number, asset tag)

7.5 Update Procedures

Application updates follow the same secure media transfer process as initial deployment: obtain, hash, transfer, verify, install, test, and update configuration management baseline documentation. No update mechanism within the application contacts external servers. Updates are entirely manual and controlled by the system administrator.

8. TEMPEST and EMSEC Considerations

This section addresses electromagnetic emanation security (EMSEC) and TEMPEST considerations for deploying Airgap Voice on commercial off-the-shelf (COTS) Apple Silicon hardware. Airgap Voice is a software application; it does not alter the hardware's electromagnetic emission profile. TEMPEST compliance is a property of the hardware and facility, not the software. However, the software's processing characteristics affect the information content of any emanations, which is the relevant concern.

8.1 Acoustic Emanation Analysis

Apple Silicon Macs present a favourable acoustic profile for classified environments:

- Many Apple Silicon Mac models operate with passive cooling (fanless). No fan means no acoustic variation correlated with processing load.
- Models with active cooling use variable-speed fans governed by thermal management firmware and do not correlate in real-time with individual application processing patterns at a granularity sufficient for information extraction.
- The Neural Engine and GPU, which perform the inference workload, do not produce user-audible acoustic emissions.
- No mechanical storage (SSD only) eliminates seek noise as an acoustic channel.

The acoustic emanation risk for Airgap Voice on Apple Silicon hardware is assessed as low. For environments requiring additional assurance, fanless models are recommended.

8.2 Electromagnetic Emission Profile

Apple Silicon integrates CPU, GPU, Neural Engine, and memory controller into a single System-on-Chip (SoC). This architecture has favourable implications for EM emissions:

- Unified memory architecture: Memory is on-package or on-substrate, reducing the length of high-speed memory traces that are a primary source of EM emanation in traditional architectures.
- No discrete GPU: Integrated GPU means no separate PCIe interface, VRAM, and power delivery — all of which are significant EM emitters in traditional workstations.
- Low power draw: Apple Silicon Macs consume significantly less power than comparable x86 systems, resulting in lower overall EM emissions.

8.3 TEMPEST Zone Deployment Guidance

- Zone A (within SCIF boundary): Deploy within the inspectable space. Standard TEMPEST countermeasures per facility accreditation apply.
- Zone B/C (controlled/uncontrolled space): Not recommended for processing classified information on COTS hardware without additional TEMPEST countermeasures (e.g., shielded enclosure, filtered power).
- Minimum countermeasures: Deploy per the facility's TEMPEST countermeasures plan. Fanless models are preferred.

8.4 Honest Limitation: No TEMPEST Certification

Standard Apple hardware is not TEMPEST-certified. No software application can change this. If the environment's TEMPEST requirements mandate the use of TEMPEST-certified equipment (per CNSSAM TEMPEST/01-13 or equivalent), then COTS Apple hardware does not meet that requirement regardless of the software installed. Mitigations include:

- Deployment within a TEMPEST-shielded enclosure or SCIF with adequate zone protection
- Use of filtered power connections
- Selection of fanless hardware models to minimise acoustic coupling
- Acceptance of residual EM risk by the AO based on threat environment assessment

9. Residual Risk Assessment

No system is risk-free. This section provides an honest assessment of the residual risks associated with deploying Airgap Voice in a classified environment, along with recommended mitigations and risk acceptance considerations.

Risk	Severity	Likelihood	Mitigation	Residual Risk Level
Physical access to endpoint allows unauthorised microphone use	High	Low (within SCIF)	Physical access controls per ICD 705; session lock; TCC permission management	Low
Model weight extraction from application bundle	Medium	Low	Inherent to any local deployment; weights are not classified data; does not enable information extraction	Low (Acceptable)
Swift runtime memory management may not zero all intermediate buffers	Medium	Medium	Critical audio buffers use memset_s for explicit zeroing; ARC may retain intermediate objects temporarily; mitigated by short processing window	Medium (Accept with POA&M)
Accessibility API observation by other privileged processes	Medium	Low	Hardened Runtime limits process injection; macOS TCC requires explicit grant; monitor authorised Accessibility clients	Low
COTS hardware not TEMPEST-certified	Variable	Variable	Deploy within TEMPEST-controlled zone; use fanless models; AO risk acceptance based on threat environment	Variable (AO Decision)
Vendor (software supply chain) compromise	High	Very Low	Code signing and notarisation provide integrity chain; SHA-256 hash verification at transfer; binary can be analysed prior to deployment	Low
macOS operating system vulnerabilities	High	Medium	Apply DISA STIG hardening; maintain OS patching via secure media transfer; this risk is not specific to Airgap Voice	Medium (Shared risk with all macOS deployments)

9.1 Risk Acceptance Recommendations

Risks acceptable without additional mitigation:

– Model weight extraction: Inherent to local deployment and does not expose classified data.

- Vendor supply chain compromise: Mitigated adequately by code signing, notarisation, and hash verification.
- Accessibility API observation: Mitigated by Hardened Runtime and TCC controls.

Risks acceptable with POA&M:

- Swift runtime memory management: Recommend POA&M item to monitor future application versions for additional memory management hardening.
- macOS OS vulnerabilities: Standard shared risk; mitigated by STIG compliance and patching.

Risks requiring AO decision based on threat environment:

- TEMPEST/EMSEC: Risk level depends on the facility's TEMPEST zone classification, threat environment, and available countermeasures. The AO must make this determination based on site-specific factors.

10. Recommendation

10.1 Summary Finding

Airgap Voice presents an architecture that is fundamentally aligned with the security requirements of classified information processing environments. Its defining characteristics — zero network activity, no persistent data storage, local-only processing, bundled models, and a licensing model that requires no activation server or license validation — address the most significant threat vectors associated with transcription software in classified environments.

The application's behaviour is empirically verifiable using standard system monitoring tools available on every macOS installation, providing the accrediting authority with concrete, reproducible evidence rather than vendor assertions alone.

> **Based on the analysis in this guide, Airgap Voice is assessed as suitable for evaluation and piloting in classified environments operating at CUI through TS/SCI, subject to the standard RMF authorisation process and the residual risk considerations identified in Section 9.**

10.2 Recommended Next Steps

- Coordinate with the Information System Security Manager (ISSM) to initiate the RMF process for adding Airgap Voice as a software component to the host system's authorisation boundary.
- Procure a site license through the appropriate acquisition channel (GSA Schedule, direct procurement, or other authorised vehicle).
- Conduct the security assessment procedures outlined in Section 6.3 using the deployed application on a representative system within the target environment.
- Document findings in the Security Assessment Report (SAR) and submit to the Authorising Official for ATO decision.
- Establish the configuration management baseline per Section 7.4.
- Implement continuous monitoring per the facility's continuous monitoring plan, incorporating the periodic verification procedures identified in this guide.