AIR GAP VOICE

Source Protection for Foreign Correspondents

April 2026

PUBLIC RELEASE

Airgap Voice is a speech-to-text application for macOS and iOS that processes all audio locally on the user's device. No audio, transcription text, or telemetry data is ever transmitted over a network. This document describes the technical architecture for security architects, compliance officers, and technical evaluators who need to understand how Airgap Voice achieves data sovereignty guarantees that cloud-based transcription services cannot provide.

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

Table of Contents

1. Introduction

The foreign correspondent occupies a critical position in journalism. Operating thousands of miles from the legal protections of their home country, often in environments where press freedom is nominal or nonexistent, the correspondent carries a responsibility that extends well beyond accurate reporting. Every source who speaks to a foreign journalist in a conflict zone is placing extraordinary trust in that journalist. The source is trusting that the correspondent will protect their identity, their words, and the very fact of the conversation. In many of the places where foreign correspondents work, that trust is a matter of life and death.

Source protection is the foundational ethical obligation of journalism. It predates digital technology by centuries. It is codified in professional ethics, supported (unevenly) by law, and understood by every working journalist as a non-negotiable commitment. When a source speaks on condition of anonymity, or when a journalist determines that identifying a source could place that person in danger, the journalist assumes an obligation that does not expire, does not yield to convenience, and does not grant exceptions.

At the same time, AI-powered transcription has become transformative for field reporting. The ability to dictate notes, transcribe interviews, and convert spoken observations into text with high accuracy and speed represents a genuine advance in the correspondent's toolkit. In environments where typing is impractical, where notes must be taken quickly, or where a journalist needs to capture their observations while they are still fresh, voice-to-text transcription is enormously valuable.

The most widely available transcription tools are cloud-based services that transmit audio to remote servers for processing. For a journalist dictating notes about a restaurant review or a city council meeting, this is unremarkable. For a journalist dictating notes about a conversation with a dissident, a whistleblower, or a witness to war crimes, it is a potential catastrophe.

This paper examines a straightforward question: can a foreign correspondent use AI transcription without compromising source safety? The answer depends entirely on how the transcription is processed, where the audio travels, and who has access to it. The distinction between cloud-based and local transcription is a matter of operational security, and for the sources who trust correspondents with their safety, it is a matter of survival.

2. The Threat Landscape for Field Journalists

2.1 State Surveillance

Governments with authoritarian tendencies invest heavily in monitoring network traffic within their borders. Deep packet inspection, keyword filtering, and bulk collection of internet metadata are standard practices in dozens of countries where foreign correspondents regularly operate. When a journalist connects to the internet from within such a country, the connection itself may be logged, analyzed, and flagged. The content of the transmission, when intercepted or compelled from the service provider, becomes available to the very authorities the journalist's sources may be exposing.

State surveillance is not limited to authoritarian regimes. Democracies also conduct signals intelligence, and the legal frameworks governing that collection vary widely. A journalist operating in one country may find that their data, transmitted to a cloud service headquartered in another country, is subject to the surveillance laws of both jurisdictions, plus any jurisdiction through which the data transits.

2.2 Hostile Non-State Actors

In conflict zones, the threat is not limited to governments. Armed groups, private intelligence services, and criminal organizations all possess varying degrees of technical capability. Some have demonstrated the ability to intercept communications, compromise devices, and exploit network vulnerabilities. A journalist working in a region controlled by such groups faces multiple advanced threats.

2.3 Border Crossings and Device Seizure

Border inspections present a distinct and well-documented risk. Customs and immigration authorities in many countries assert the right to inspect electronic devices, including the authority to compel the traveler to unlock the device and provide access to its contents. For a journalist carrying notes from sensitive interviews, a border inspection can expose source material directly. When a journalist has used an online transcription service, the device may contain not only the cached transcribed text but also evidence of the cloud connection: account credentials, API keys, usage logs, and cached audio files that reveal the scope and nature of the journalist's work.

2.4 Compromised Network Infrastructure

In conflict zones and countries with weak rule of law, network infrastructure cannot be trusted. Internet service providers may be state-controlled or compromised by hostile actors. Man-in-the-middle attacks on local networks are documented and recurring. Hotel Wi-Fi networks and press center internet connections are notoriously insecure and are frequently monitored by intelligence services, both domestic and foreign. A journalist who connects to any network in such an environment must assume that the connection is observed.

2.5 The Metadata Problem

Even when the content of a transmission is encrypted, metadata tells its own story. The fact that a journalist connected to a transcription service at a particular time, from a particular location, for a particular duration, generates a pattern of activity that can be correlated with other information. If the journalist interviewed a source at 2:00 PM and connected to a cloud transcription service at 2:45 PM, the metadata alone narrows the field of inquiry for anyone seeking to identify the source. Metadata is famously difficult to suppress.

Cloud services generate metadata on their own servers as well: timestamps of API calls, durations of audio submissions, IP addresses, device identifiers, and account activity logs. This metadata exists outside the journalist's control, on infrastructure the journalist cannot inspect, subject to legal processes the journalist may never learn about.

2.6 Legal Compulsion

Vendors who receive and process audio are subject to the legal systems of the jurisdictions in which they operate. Court orders, subpoenas, national security letters, and mutual legal assistance treaty requests can compel a vendor to produce records, including audio files, transcripts, and metadata. In many jurisdictions, the vendor may be prohibited from notifying the user that such a request has been made. The journalist's privilege, where it exists, protects the journalist from being compelled to testify; it does not extend to third-party vendors who possess copies of the journalist's source material.

3. Cloud Transcription as an Operational Security Risk

When a journalist dictates notes into a cloud transcription service, the following sequence occurs: the journalist's voice is captured by the device microphone, converted to a digital audio stream, transmitted across a network to a remote server, processed on that server, and the resulting text is transmitted back. At every stage of this process, the source's identity, statements, and context are exposed to risk.

The audio traverses network infrastructure that may be monitored by the very actors the source is exposing. It arrives at a server operated by a third-party vendor, where it is processed in cleartext (the server must access the unencrypted audio to perform the transcription). The vendor's employees, the vendor's subcontractors, and the vendor's automated systems all have potential access to the audio and the resulting transcript. The data may be stored, logged, cached, or used for model improvement, depending on the vendor's policies and practices.

3.1 The Vendor as a Compellable Third Party

The moment audio leaves the journalist's device and arrives on a vendor's server, it exists in a location that the journalist does not control and cannot protect. The vendor is subject to legal process in its own jurisdiction. If a government demands the transcription records of a particular user, the vendor must comply or litigate, and the journalist may never be informed. The "third-party doctrine," established in multiple legal systems, holds that information voluntarily shared with a third party carries reduced privacy protections. By using a cloud service, the journalist has voluntarily shared the source material with the vendor.

3.2 The Chilling Effect

Sources are not naive about technology. A whistleblower inside a government ministry, a defector from an armed group, or a witness to state violence understands the risks of digital communication. When a source learns that the journalist's transcription tool sends audio to a cloud server, the source's willingness to speak may evaporate. The chilling effect is real and measurable: sources who believe their words will be processed by third-party systems are less likely to speak, less forthcoming when they do, and more likely to impose conditions that limit the journalist's ability to collect further critical information.

3.3 Historical Patterns

The history of journalism in the digital age includes a troubling pattern of source exposure through technology. Email records, phone metadata, messaging app data, and cloud storage contents have all been used to identify and target journalists' sources. In multiple documented cases, governments have obtained records from technology companies to identify individuals who communicated with journalists. The lesson is consistent: any system that creates a record of source communication on a third-party server is a potential vector for source identification.

3.4 The False Comfort of Encryption in Transit

Cloud transcription vendors typically encrypt data in transit using TLS. This protects against casual interception on the network path, but it does not address the fundamental problem. The encryption terminates at the vendor's server, where the audio must be decrypted for processing. "End-to-end encryption" is a term that is frequently misapplied to cloud services; in a cloud transcription context, one "end" is the journalist's device and the other "end" is the vendor's server. The vendor is the endpoint, and the vendor has full access to the cleartext audio. For a journalist concerned about source protection, encryption in transit solves the wrong problem.

4. Source Protection Principles Applied to Transcription

4.1 The Journalist's Privilege and Its Limits

The journalist's privilege, also known as the reporter's privilege, provides varying degrees of legal protection for journalists who refuse to disclose their sources. In the United States, most states have shield laws, but there is no federal shield law, and the privilege has been narrowly construed by federal courts. In Europe, the European Court of Human Rights has recognized the importance of source protection, but national implementations vary widely. In many countries where foreign correspondents work, the privilege does not exist at all.

Critically, the journalist's privilege protects the journalist from being compelled to testify or produce records. It does not protect third parties who happen to possess copies of the journalist's source material. A cloud transcription vendor that holds audio of a journalist's dictated notes is not covered by the journalist's privilege. The vendor can be subpoenaed, and the journalist may have no standing to challenge the subpoena.

4.2 The Reasonable Precautions Standard

Courts and professional bodies evaluate source protection claims in part by assessing whether the journalist took reasonable precautions to protect the source's identity. A journalist who transmits source material to a cloud service without considering the operational security implications may find that a court views this as a failure to take reasonable precautions.

4.3 Minimizing the Attack Surface

In operational security, the attack surface is the sum of all points where an adversary can attempt to access protected information. Every system that touches source material, every network the data traverses, every server where it is stored or processed, represents an additional point of vulnerability. The principle is straightforward: the fewer systems that touch source material, the fewer opportunities for compromise. A transcription workflow that keeps audio on the journalist's device has an attack surface of one. A workflow that transmits audio to a cloud server has an attack surface that includes the device, the network path, the vendor's server infrastructure, the vendor's employees, and the vendor's compliance with legal demands.

4.4 Data Minimization

The principle of data minimization holds that sensitive information should exist in as few places as possible, for as short a time as necessary. This principle is embedded in data protection regulations worldwide and is a cornerstone of operational security practice. Applied to transcription, it means that source-related audio should not be unless there is a compelling reason to do so, and that any copies should be deleted as soon as they are no longer needed. Cloud transcription services create copies that the journalist cannot verify have been deleted, on systems the journalist cannot audit.

4.5 Professional Codes and Standards

The Society of Professional Journalists' Code of Ethics calls on journalists to "recognize that gathering and reporting information may cause harm or discomfort" and to "recognize that legal access to information differs from an ethical justification to publish." The International Federation of Journalists' Declaration of Principles and the Reporters Without Borders recommendations on digital security both emphasize the obligation to protect sources through all available means. These codes do not yet specifically address transcription tools, but the principles they articulate apply directly: a journalist who has the option to process source material locally, without network transmission, and chooses instead to transmit it to a cloud service, is not exercising the level of care these codes envision.

5. Local Transcription as Operational Security

Local transcription, in which audio is processed entirely on the journalist's own device without any network transmission, addresses the threat landscape by eliminating the vectors rather than attempting to mitigate them. This is an important distinction. Network encryption mitigates the risk of interception; local processing eliminates it. Vendor privacy policies mitigate the risk of data misuse; the absence of a vendor eliminates these risks.

5.1 Eliminating the Network Vector

When audio is processed on the device, it never traverses a network. There is nothing to intercept, nothing to log at an ISP, nothing to capture via a man-in-the-middle attack on a compromised hotel network. The network vector, which represents the largest and most diverse category of threats to cloud-based transcription, simply does not exist. Tools like Airgap Voice process speech-to-text entirely on the journalist's Mac or iPhone, with zero network transmission. The journalist can verify this independently using standard network monitoring tools.

5.2 No Server, No Subpoena Target

If there is no server, there is nothing to subpoena. No court order can compel a vendor to produce records that do not exist. No national security letter can demand logs from a server that does not exists. The third-party doctrine does not apply, because there is no third party. The journalist retains full control of the only copy of the audio and the transcription, and the journalist's privilege (where it exists) applies to materials in the journalist's possession without the complication of data copies on third party systems.

5.3 No Account, No Metadata Trail

A local transcription tool that requires no account, no login, and no internet activation creates no metadata trail linking the journalist to a transcription service. There is no user record, no usage history, no connection log. At a border crossing, there are no API keys, no cloud service credentials, and no account history on the device that would reveal the journalist's use of transcription tools. The tool is simply an application installed on the device, indistinguishable from dozens of other productivity applications.

5.4 Offline Capability

In many operational environments, connecting to a network is itself a risk. The act of connecting generates metadata, exposes the device to network-based attacks, and may trigger surveillance systems. A transcription tool that works fully offline, with no network connection required, allows the journalist to transcribe source material without ever exposing the device to a network. The lesson from the war in Ukraine is direct and well documented. On both sides of the conflict, military positions were identified and struck because a single service member connected a personal device to a cellular network. The opposing force used signal intelligence, geolocation data, and open-source analysis to extract location coordinates from these digital emissions, and artillery or drone strikes followed within minutes. A not insignificant number of soldiers were killed because of a careless digital fingerprint. The same risks apply to journalists reporting from a conflict zone where network traffic is actively monitored, for example in Gaza, where any digital devices become targets because of a single transmission announces its location.

The operational answer is straightforward: disable all wireless radios, place the device in a Faraday enclosure, and use tools that continue to function without any network connection. A local transcription system that operates entirely offline allows the correspondent to continue working, recording observations, transcribing interviews, and drafting reports, without producing the network emissions that could make them a target.

A more common risk is in jurisdictions where internet traffic is actively monitored, that journalist information is silently intercepted, analyzed, and then used to establish ground to deny visa renewal or re-entry to the country in case the reporting is seen unfavorable. The plausible deniability of this practice makes it particularly hard to prove, but the consequence for foreign correspondence is very real. The operational answer, similarly, rests on offline tools that cannot be intercepted until a safe means of data transmission has been establishes.

5.5 Multi-Language Support

Foreign correspondents work across language boundaries on a regular base. A correspondent covering a story that spans multiple countries may need to transcribe notes in several languages within a single day. A local transcription tool that supports 52 languages and dialects allows the correspondent to use a single tool across all their working languages. Each additional service account is an additional point of exposure, an additional set of credentials on the device, and an additional vendor who holds records of the journalist's activity.

5.6 Verifiability and Auditability

A vendor's privacy policy is a statement of intent, not a technical guarantee. Local transcription offers something that cloud services cannot: independent verification. A journalist can confirm that their transcription tool is not transmitting data by using standard network monitoring tools available on any Mac or iPhone. A newsroom security team can audit the tool's network behavior independently. This verifiability is important not only for the journalist's own confidence but also for building trust with sources. A journalist who can demonstrate to a source that the transcription tool operates entirely offline is offering a concrete, verifiable assurance.

6. Field Deployment Considerations

6.1 Pre-Deployment Preparation

The time to install and configure transcription tools is before entering the operational area, while on a trusted network. Download the application, and verify correct operation while connectivity is not a constraint. Once in the field, the tool operates without any network.

6.2 Air-Gapped Operation

For transcription sessions involving sensitive sources, the strongest posture is to disable Wi-Fi and cellular data before beginning. With a local transcription tool, this has no effect on functionality; the tool operates identically whether the device is connected or not. The journalist can verify that the device is fully disconnected before dictating, providing an additional layer of assurance that no data can leave the device during the session.

6.3 Device Hygiene and Secure Deletion

Local processing means that the audio exists only on the journalist's device during the transcription process. After transcription, the audio data in volatile memory is released. The transcribed text, however, does exist in whatever application the journalist dictated into. Standard practices for secure handling of sensitive notes apply: use encrypted storage, follow established protocols for secure deletion when notes are no longer needed, and maintain the same discipline with transcribed text as with any other source material.

6.4 Multi-Language Workflows

A single tool that handles multiple correspondent's working languages reduces the number of applications that touch sensitive material. Every additional application in the workflow is an additional point of potential failure, an additional set of permissions to manage, and an additional tool whose security properties must be evaluated. Consolidation reduces complexity, and in operational security, complexity is the enemy.

6.5 Integration with Existing Workflows

Local transcription tools that deliver text directly to the active application, at the cursor position, allow the journalist to dictate into their encrypted notes application, their word processor, or their messaging client without additional steps. The system clipboard is bypassed for security reasons because on modern Mac OS, the operation system maintains a history of the clipboard. Because AirGap Voice does not use the clipboard, there is no hidden copy.

There is no separate transcription application to manage, no copy-and-paste operation, and no intermediate file to secure. The text appears where the journalist needs it, in the application the journalist already knows and trusts.

6.6 Travel Across Borders

A journalist whose transcription tool requires no cloud service connection carries no evidence of cloud transcription usage on their device. There are no cloud service accounts to discover during a device inspection, no API keys to raise questions, no usage logs that reveal the timing and volume of transcription activity. The tool presents as a standard productivity application. Combined with device encryption and strong access controls, this profile minimizes the information available to border authorities during a device inspection.

6.7 Newsroom Deployment

For news organizations deploying transcription tools to correspondents in the field, local processing eliminates significant operational overhead. There is no cloud infrastructure to manage, no accounts to provision or deprovision, no vendor relationship that creates a discoverable link between the news organization and the transcription of sensitive material. Each correspondent's installation is independent, self-contained, and leaves no organizational footprint on external systems.

7. Recommendations for News Organizations

News organizations that deploy correspondents to high-risk environments should consider transcription security as an integral part of their operational security framework. The following recommendations are offered as a starting point for organizations seeking to align their transcription practices with their source protection obligations.

- Conduct a transcription audit. Identify every tool currently used by staff for dictation and transcription. For each tool, determine where the audio is processed, who has access to it, and what records the vendor retains. Many organizations will find that correspondents are using personal accounts with consumer cloud services, creating unmanaged exposure.
- Establish a transcription policy for sensitive source material. At minimum, the policy should require local-only processing for any dictation involving sources, whistleblowers, or sensitive subjects. Cloud transcription may be acceptable for non-sensitive material such as general notes, non-confidential interviews, and administrative dictation.
- Include transcription tools in pre-deployment security briefings. Correspondents preparing for assignments in high-risk environments should receive specific guidance on transcription security, including which tools are approved for sensitive material and how to verify that a tool is not transmitting data.
- Train correspondents on verification. Every correspondent should know how to confirm that their transcription tool is not making network connections. On macOS, this can be verified using Activity Monitor or the "nettop" command-line tool. On iOS, network activity indicators and device settings provide confirmation. This training should be practical, with hands-on exercises rather than written instructions alone.
- Integrate transcription security into existing digital security protocols. Most news organizations with correspondents in high-risk environments already have protocols covering encrypted messaging, VPN usage, device encryption, and secure file transfer. Transcription should be added to this framework with the same level of rigor.
- Consider the transcription tool as part of the source protection kit. Alongside Signal for messaging, Tor for anonymous browsing, and encrypted storage for files, a local transcription tool should be part of the standard suite of tools deployed to correspondents who handle sensitive source material. Source protection is not a single tool; it is a practice supported by a set of tools, each addressing a specific vector.
- Review vendor claims critically. For any transcription tool considered for deployment, verify the vendor's claims about data handling independently. A tool that claims to process locally but requires an internet connection to function is not truly local. A tool that claims not to collect data but requires an account is creating a record. Independent technical verification should be a prerequisite for approval.

8. Conclusion

Source protection is a commitment to be honored. It predates the internet, predates digital recording, and predates every technology that journalists currently use. It is rooted in the recognition that people who share sensitive information with journalists, particularly in dangerous environments, are placing their safety in the journalist's hands. No productivity gain, no matter how significant, justifies compromising that trust.

The most secure transmission is the one that never happens. When audio is processed on the journalist's own device, with no network connection, no cloud server, no vendor, and no third-party access, the transmission vector is eliminated entirely. This is not a risk mitigation; it is a risk elimination. The difference matters, because the consequences of failure are measured in human safety.

Cloud transcription asks journalists to trust a vendor's privacy policy, a vendor's security practices, and a vendor's willingness to resist legal compulsion. Local transcription eliminates the need for that trust by eliminating the vendor's access to the material. The journalist does not need to evaluate the vendor's privacy policy, because the vendor never receives any data. The journalist does not need to worry about subpoenas to the vendor, because the vendor holds no records.

Foreign correspondents can now use AI transcription with the same confidence they place in a notebook and pen. The technology has matured to the point where high-quality, multi-language speech recognition can run entirely on a journalist's Mac or iPhone. The words stay with the journalist, on the journalist's device, under the journalist's control.

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