



Audio Intelligence Privacy Overview

September 2026

Contents

- Introduction 3**
- Privacy and security architecture 4**
 - Secure Exclave 4
 - End-to-end encrypted syncing with iCloud 5
- Siri Recap 6**
 - Detecting that a conversation is happening 6
 - Audio securely transfers to the Secure Exclave on iPhone 6
 - Secure on-device transcription turns audio into text 6
 - Private Cloud Compute summarizes the text 6
 - Your summary in the Siri app 7
 - How Siri Recap respects those around you 7
 - Control for Siri Recap.....7
- Live Rewind.....9**
 - A continuous rolling buffer.....9
 - Secure on-device processing turns audio into text 9
 - Review and optionally save.....9
 - How Live Rewind respects those around you.....9
 - Control for Live Rewind 9
- Shazam and Sound Recognition.....10**
- Conclusion 11**

Introduction

Audio Intelligence brings a new class of Apple Intelligence capabilities to Apple Watch. It is designed to help you stay present and engaged, catch details in a conversation you may have missed, identify music playing nearby, and be aware of important sounds around you, all while protecting your privacy and the privacy of those around you.

Audio Intelligence on Apple Watch Series 12 and Apple Watch Ultra 4 relies on the tight integration of hardware and software with the latest AI models to provide four new features.



Sound Recognition on Apple Watch detects important sounds like sirens, alarms, and doorbells on-device.

Siri Recap

Siri Recap creates high-level Apple Intelligence-generated summaries of your conversations throughout your day so you can stay present and engaged in the moment.

Live Rewind

Live Rewind helps you recall what was just said in conversation. A double press of the Digital Crown shows you the previous 15 seconds of a conversation as text, so you can catch something you may have missed or are less familiar with.

Music Recognition with Shazam

Shazam instantly detects the music playing around you and automatically displays the song title and artist's name on the Music Recognition widget in the Smart Stack, all without a tap.

Sound Recognition on Apple Watch

Sound Recognition on Apple Watch detects important sounds like sirens, alarms, and doorbells, and it notifies you of those sounds, even when your iPhone is not with you. This is especially helpful for people who are deaf or hard of hearing.

All four of these features are enabled by the new S11 chip on Apple Watch and its Secure Enclave, a dedicated secure buffer that processes audio privately in a protected environment that no one can access, not even Apple.

Audio Intelligence features do not create audio recordings that can be accessed by the operating system, apps, the user, or Apple. There is no recording to share, forward, or produce if requested by any party, because no recording exists.

This paper details how Audio Intelligence works, the privacy and security architecture that protects you and those around you, and the controls you have over the entire experience.


Privacy and security architecture

Secure Exclave

On Apple Watch

The S11 chip in Apple Watch Series 12 and Apple Watch Ultra 4 includes a Secure Exclave: a hardware-isolated compartment built into the silicon to process sensor data separately from the rest of the system. Data that is processed and analyzed in the Secure Exclave cannot be accessed by watchOS, apps, the user, or Apple.

Audio from the microphone enters the Secure Exclave of Apple Watch, where it is initially processed for speech, sounds, or music, without transcribing or storing the raw audio. This buffer is a continuously overwritten stream that exists only within the protected hardware and never creates an audio recording.



Audio Intelligence is only possible because of the Secure Exclave on our new S11 chip. The Secure Exclave processes raw audio in a contained, secure environment.

On iPhone

iPhone 16, iPhone 16 Pro, iPhone 17e, iPhone Air, and later models contain a Secure Exclave that processes raw audio in a contained, secure environment that cannot be accessed by iOS, apps, the user, or Apple.

Secure pairing between Apple Watch and iPhone

Apple Watch pairs with iPhone using a new audio-verified pairing mechanism that exists in addition to Bluetooth pairing. This creates an even more secure connection between the Secure Exclaves on Apple Watch and iPhone. Secure pairing ensures that ambient data from the Secure Exclave on Apple Watch can only be sent to the Secure Exclave of the paired iPhone, and no other device.

Encrypted communication between Secure Exclaves

For Siri Recap, the Secure Exclaves on Apple Watch and iPhone establish a secure, encrypted channel during the secure pairing process. Then, raw audio is encrypted and transferred from the Secure Exclave on Apple Watch to the Secure Exclave on the paired iPhone. Together, these two hardware-protected spaces form a direct secure chain where encrypted raw audio moves between them and is inaccessible to the operating systems, apps, the user, or Apple.

If Apple Watch or iPhone is lost

The encryption keys on the Secure Exclaves are device-bound and time-bound. The keys rotate regularly and expire automatically. If your Apple Watch or iPhone is lost, you can also remotely wipe it using Find My, which also invalidates the pairing keys in the Secure Exclaves immediately.

End-to-end encrypted syncing with iCloud

Once a Siri Recap or Live Rewind text is generated, you decide what to keep. No raw audio is ever saved or synced. If you use iCloud with two-factor authentication and have a device passcode, any Audio Intelligence text you choose to save syncs across your devices end-to-end encrypted. Only you can access this information, and only on devices where you are signed in to iCloud. Apple does not hold the encryption keys. No one else, not even Apple, can access the end-to-end encrypted data you saved in the Siri app. If you delete a saved Siri Recap or Live Rewind text, it is permanently removed from all your devices and from iCloud.

Siri Recap


Siri Recap summarizes conversations throughout your day and produces high-level Apple Intelligence-generated notes so you can stay present in the moment and catch up later.

Detecting that a conversation is happening

The S11 chip uses a lightweight AI model to determine whether speech is occurring nearby. This model doesn't transcribe or record anything. It simply detects whether a conversation has started. If a conversation is detected, audio flows into a protected buffer inside the Secure Enclave on the S11 chip. It is isolated in hardware so that it is inaccessible to the user, the operating system, apps, and even Apple.

Audio securely transfers to the Secure Enclave on iPhone

The Secure Enclave on Apple Watch encrypts the audio and transmits it to the Secure Enclave on the paired iPhone. Once the transfer completes, audio is immediately deleted from the Secure Enclave on Apple Watch, and only the Secure Enclave of the paired iPhone can decrypt it. If the transfer fails, for example if iPhone is not within wireless range of Apple Watch, the Secure Enclave on Apple Watch will attempt the transfer again when the connection is re-established. If a connection cannot be re-established in time, the encryption keys expire and the encrypted audio is permanently deleted.



Siri Recaps will have an Apple Intelligence generated title, summary, and key points. You can review them in the Siri app.

Secure on-device transcription turns audio into text

The encrypted audio is decrypted inside the Secure Enclave of the paired iPhone. On-device speech recognition transcribes the audio to text, and then an on-device language model generates a condensed version that is less than half the length of the original transcript. This is designed to distill the conversation to its core meaning, removing content suggestive of conversation tone, non-essential language, filler words, and redundancies, while preserving the topics discussed. Once processed, the raw audio is immediately and permanently deleted and no longer exists on either Secure Enclave. Only the condensed text remains. Finally, a safety model screens the text to omit potentially harmful terms.

Private Cloud Compute summarizes the text

The condensed text is encrypted and sent to Private Cloud Compute. Contextual information is also sent to improve summary quality. For example, Now Playing data helps the model understand that speech may have originated from music or a podcast. Calendar data can help generate accurate titles of summaries. High-level location labels such as home, work, and school, along with locality information such as city, state, and country, add context to the summarization. Point-of-interest categories like "grocery store" or "park" may also be included. Precise location and specific points-of-interest are not included.


Apple Foundation Models running on Private Cloud Compute generate a title and a summary with key points. Siri Recap is designed so that sensitive information is not included in summaries, such as financial data, government-assigned identifiers, authentication data, and certain personal identifiers.

Data in Private Cloud Compute is never stored or accessible by Apple or anyone else. It is used only to fulfill your request. Independent security experts can verify this promise at any time. For a complete description of Private Cloud Compute's architecture and guarantees, see Apple's security documentation at security.apple.com/blog/private-cloud-compute/.

Your summary in the Siri app

The finished, text-based summary is encrypted and sent from Private Cloud Compute back to iPhone and Apple Watch, where it is available to view in the Siri app in the Recaps tab.

Siri Recaps are automatically deleted after 7 days, unless you explicitly choose to save them. If you use iCloud with two-factor authentication and have a device passcode, Siri Recaps sync across your devices end-to-end encrypted. You can edit or delete any summary directly in the Siri app.



With Siri Recap, you can control when and where it is active, or control it manually.

How Siri Recap respects those around you

By design, Siri Recap does not create a recording, does not produce a verbatim transcript, and does not identify and attribute speakers. The output is a brief, high-level summary, comparable to notes a person might write after a conversation. There is no audible signal because no raw audio is retained, and there is no way to reconstruct the original audio from a Siri Recap or share raw audio with anyone.

Control for Siri Recap

To use Siri Recap, you must opt in when setting up Apple Watch, and you can enable or disable Siri Recap at any time through Siri Settings on Apple Watch, and Siri Settings in the Apple Watch app on iPhone. You can control when it's active (for example, only during work hours, or never at night) and where it's active (for example, only at your workplace). You can also turn it on and off manually with Control Center.

Your data's journey for Siri Recap

① Audio stream during speech detection

Where it's processed: S11 Always On Processor

Accessible to: No one. Raw audio is processed on-device and discarded in real time.



② Buffered audio on Apple Watch

Where it's processed: Secure Exclave on Apple Watch

Accessible to: No one. Raw audio is isolated in secure compartment.



③ Buffered audio on iPhone

Where it's processed: Secure Exclave on iPhone, after being transmitted encrypted from Apple Watch

Accessible to: No one. Raw audio is isolated in secure compartment.



④ Condensed text on iPhone

Where it's processed: Condensed text is encrypted and transferred from Secure Exclave on iPhone to Private Cloud Compute. If you have optionally enabled your Apple Intelligence Report in Privacy & Security Settings on your iPhone, you can verify the Private Cloud Compute request in that report and only on your device. Depending on your settings, it can be visible to you for 15 minutes or 7 days.

Accessible to: Only you. Your optional Apple Intelligence Report on your device is only accessible to you after you first authenticate to view it.



⑤ Condensed text on Private Cloud Compute

Where it's processed: Condensed text is decrypted on Private Cloud Compute once it is received, where it is processed (subject to daily usage limits).

Accessible to: No one. Condensed text in Private Cloud Compute cannot be accessed by anyone, not even Apple. Private Cloud Compute creates a summarization from the condensed text and removes sensitive information. The summarization is then encrypted and sent back to iPhone. The results are immediately deleted from Private Cloud Compute.



⑥ Summarization

Where it's processed: Siri app, and synced across your devices end-to-end encrypted via iCloud.

Accessible to: Only you.



Siri Recap can be turned on and off at anytime right from Control Center.

Live Rewind

Live Rewind lets you instantly see the last 15 seconds of a conversation as text. This is great for moments like when someone quickly shares a book recommendation, cooking instructions, or even an unfamiliar term you want to revisit later.

A continuous rolling buffer

Audio from the microphone flows into the audio buffer inside the Secure Exclave on the S11 chip. Old audio is continuously replaced by new audio, so audio doesn't accumulate, and raw audio isn't accessible outside this protected buffer.

Secure on-device processing turns audio into text

Live Rewind processes audio on-device. Live Rewind requires a double-press of the Digital Crown to activate it every time you use it. After activating, the immediately preceding 15 seconds of buffered audio is sent from the Secure Exclave of the Apple Watch to iPhone. If iPhone is not within wireless range of Apple Watch, the transfer will fail with an error, and the audio will be immediately discarded. If the transfer to iPhone is successful, on-device speech recognition converts the audio to text, and is designed so that the raw audio cannot be accessed by iOS, apps, the user, or Apple. Once converted, the raw audio is then discarded immediately.

Review and optionally save

The 15-second Live Rewind text is then sent back to Apple Watch, where it appears on the display. 30 seconds after your Apple Watch display dims, the text disappears. You can choose to save it to the Siri app to revisit later, where it syncs end-to-end encrypted across your devices. You can also ask Siri about the content of the Live Rewind text, at which point only the text is sent to Private Cloud Compute for processing.

How Live Rewind respects those around you

When activated, an audible chime plays from the speaker on your Apple Watch, even if your Apple Watch is on silent or you have headphones connected. A full-display animation and microphone indicator appear on the Apple Watch display, and a distinct double-press gesture is required to activate Live Rewind. This is to alert those nearby that the feature has been activated.

Control for Live Rewind

You can opt in to Live Rewind when setting up Apple Watch, or you can enable and disable it at any time through Siri Settings on Apple Watch, and Siri Settings in the Apple Watch app on iPhone. Every time you use Live Rewind, it requires a double-press of the Digital Crown.



Live Rewind help users recall what was just said in the last 15 seconds of a conversation. You can ask Siri about the content or save it to the Siri app to revisit later.

Shazam and Sound Recognition

The Secure Exclave in the S11 chip also enables Shazam and Sound Recognition on Apple Watch. Audio is buffered into small segments within the Secure Exclave and analyzed for specific sounds, then immediately deleted if nothing is detected in that buffer.

When Shazam is enabled, the Secure Exclave on Apple Watch is designed to detect music playing around you. If you open the Smart Stack on Apple Watch and music has been detected, the Secure Exclave will process the audio in the secure buffer and create an acoustic signature of that song (a compact representation that cannot be reversed into the original audio). That signature, and not the raw audio, is then sent to Shazam servers to identify the song. The audio buffer on the Secure Exclave of Apple Watch is then immediately deleted. If there's a match to the signature, the song title and artist's name are sent to Apple Watch, which then displays this information in the Music Recognition widget in the Smart Stack, the Shazam app, or as a response from Siri.

For Sound Recognition, when the audio processing within the Secure Exclave recognizes a specific sound, such as a smoke alarm, doorbell, or baby crying, it sends you a notification. You can enable or disable individual sound categories in Settings. No audio is accessible outside of the Secure Exclave on Apple Watch, and no audio is transcribed or transmitted. Sound Recognition takes place on Apple Watch.



Shazam can instantly detect the music playing around you and automatically display the song title and artist's name in the Music Recognition widget.

Conclusion

Audio Intelligence features on Apple Watch help you stay present and engaged, catch details in a conversation you may have missed, and be aware of important sounds around you. These are enabled by a privacy and security architecture in which raw audio is handled within dedicated hardware, is not saved as a file, and is not accessible to the operating system, apps, or Apple. All of this is possible due to the Secure Exclave architecture on the new S11 chip and on supported iPhone models. This architecture creates a hardware-enforced boundary around audio processing. You control whether and when these features are active.

Audio Intelligence features on Apple Watch respect you by having privacy designed into the product, from the silicon up, with hardware and software working seamlessly together.

To learn more about Apple's commitment to privacy, go to apple.com/privacy.

© 2026 Apple Inc. All rights reserved. Apple and the Apple logo, Apple Watch, Apple Watch Ultra, Find My, Shazam, Siri, and watchOS are trademarks of Apple Inc., registered in the U.S. and other countries. Apple Intelligence is a trademark of Apple Inc. iCloud is a service mark of Apple Inc., registered in the U.S. and other countries. iOS is a trademark or registered trademark of Cisco in the U.S. and other countries and is used under license. Other product and company names mentioned herein may be trademarks of their respective companies. Product specifications are subject to change without notice. This material is provided for information purposes only; Apple assumes no liability related to its use. Not all features are available everywhere. See www.apple.com/ios/feature-availability and www.apple.com/watchos/feature-availability for feature availability in your region or language. September 2026