



June 28, 2019

The Honorable Christopher A. Coons
United States Senate
218 Russell Senate Office Building
Washington, D.C. 20510

Dear Senator Coons,

Thank you for your May 23, 2019 letter regarding Amazon's privacy and data security practices with respect to our Alexa voice service.

At Amazon, customer trust is our highest priority, and we know we must get privacy right in order to meet our customers' high expectations. Our customer-centric approach has led us to follow privacy-by-design principles since we introduced the Echo device and Alexa service. For example, as we described in our response to your letter from last year, we designed Echo devices and Alexa to use on-device "keyword spotting" technology to detect when a customer intends to interact with Alexa; to use visual and audible signals to clearly indicate to customers when audio is being recorded for streaming to the cloud; to continually attempt to determine when a customer's request to Alexa has ended so we can minimize the amount of audio we stream to the cloud; to allow customers to see, hear, and delete the audio that was streamed to the cloud; and to let customers control when their Echo device's microphone is enabled through a microphone on/off button. We use the customer data we collect to provide the Alexa service and improve the customer experience, and our customers know that their personal information is safe with us.

The answers to your questions are as follows:

1(a). How long does Amazon store the transcripts of user voice recordings?

We retain customers' voice recordings and transcripts until the customer chooses to delete them.

1(b). Do users have the ability to delete any or all of these transcripts?

Customers can review, listen to, and delete voice recordings associated with their account using the Voice History feature available in the Alexa app and the Alexa Privacy Hub, located at www.amazon.com/alexaprivacy. Customers can delete individual voice recordings, voice recordings from particular timeframes, or all of their voice recordings.

When a customer deletes a voice recording, we delete the transcripts associated with the customer's account of both of the customer's request and Alexa's response. We already delete those transcripts from all of Alexa's primary storage systems, and we have an ongoing effort to ensure those transcripts do not remain in any of Alexa's other storage systems. We do not store the audio of Alexa's response. However, we may still retain other records of customers' Alexa interactions, including records of actions Alexa took in response to the customer's request. And when a customer interacts with an Alexa skill, that skill developer may also retain records of the interaction. For example, for many types of Alexa requests – such as when a customer subscribes to Amazon Music Unlimited, places an Amazon Fresh order, requests a car from Uber or Lyft, orders a pizza from Domino's, or makes an in-skill purchase of premium digital content – Amazon and/or the applicable skill developer obviously need to keep a record of the transaction. And for other types of Alexa requests – for instance, setting a recurring alarm, asking Alexa to remind you of your anniversary, placing a meeting on your calendar, sending a message to a friend – customers would not want or expect deletion of the voice recording to delete the underlying data or prevent Alexa from performing the requested task.

1(c). Are there any transcripts that a user cannot delete? If so, why is a user permitted to delete voice recordings but not the corresponding text transcript?

Please see the answer to 1(b).

1(d). For what purpose does Amazon use these transcripts?

Alexa is designed to get smarter every day – this is accomplished through Amazon’s cloud services, including machine learning software. Our speech recognition and natural language understanding systems use machine learning to adapt to customers’ speech patterns and vocabulary, informed by the way customers use Alexa in the real world.

To work well, machine learning systems need to be trained using real world data. Speech is nuanced, dynamic, and has tremendous variation based on region, dialect, context, environment, and the individual speaker, including their age, whether they are a native or non-native speaker of the language, and whether they have a speech impairment. Training Alexa with voice recordings and transcripts from a diverse range of customers helps ensure Alexa works well for everyone.

In addition to using the transcripts to improve Alexa and the customer experience, we use the transcripts to provide transparency to our customer about what Alexa thought it heard and what Alexa provided as a response. Our Alexa’s Voice History feature allows customers to play the actual audio that was streamed to the cloud, review the text transcript of what Alexa thought the customer said, and review Alexa’s response. This helps customers to understand how Alexa works.

If Alexa makes a mistake, this helps customers to understand what may have gone wrong (for instance, customers can see and hear how Alexa may have misunderstood a particular word or phrase). Providing customers with the transcript also allows customers to understand and inspect exactly what Alexa is, and is not, recording.

1(e). Does Amazon take any measures to anonymize user identity and other information in relation to these transcripts?

As described in answer 1(d), we allow customers to review the transcripts of their Alexa interactions as part of our Voice History feature, which requires us to keep the transcripts associated with the customer’s account.

2. Does Amazon preserve, either in audio or textual form, the Alexa system’s responses to user questions? If so, is a user able to delete the audio or the transcript of such a response?

Please see the answer to 1(b).

3(a). In determining when a user has finished issuing a command, for how long does the system wait until it stops recording, on average?

Once activated by the wake word or the action button, the Echo device opens an audio stream to the cloud and sends the request to Alexa to respond accordingly. The audio stream closes immediately once our automatic speech recognition system determines the customer has stopped speaking the request. A blue light illuminates on the Echo device to indicate when audio is being streamed to the cloud, and customers can also enable an audible tone that plays when their Echo device begins and ends streaming audio to the cloud.

3(b). Is any audio that is captured by the on-device buffer ever sent to the cloud if the wake word is not detected?

No audio is sent to the cloud unless the Echo device detects the wake word (or customers press a button to speak to Alexa).¹

3(c). Is Audio that is captured and stored in the device's temporary memory transcribed by the automatic speech recognition system? If so, are those transcripts similarly stored only in the device's temporary memory, or are they sent to the cloud?

The short buffer on the device is continuously overwritten and is not transcribed by the automatic speech recognition system. Instead, our wake word technology works by identifying acoustic patterns that match the wake word. Once the Echo device has detected the wake word, a recording is sent to our cloud where our systems for speech recognition and natural language understanding process and respond to the customer's request. Alexa is designed to record and process as little audio from customers as possible. Processing audio that was not intended for Alexa is costly, provides no value to Amazon, and is detrimental to our customers' Alexa experience.

3(d). Amazon's July 17, 2018 letter indicates that the Alexa system comes with a setting whereby a user can allow Alexa to respond to a series of requests without the customer needing to repeat the wake word.

(i) Is this a default setting, or does a consumer need to affirmatively enable this setting?

"Follow-Up Mode" allows customers to ask Alexa multiple questions and commands without having to use the wake word each time. Follow-Up Mode must be affirmatively enabled by the customer in their Alexa app. Like other customer requests to Alexa, when a customer has enabled Follow-Up Mode, Alexa will end the stream immediately once our automatic speech recognition system determines the customer has stopped speaking to Alexa. A blue light illuminates on the Echo device to indicate when audio is being streamed to the cloud, and customers can also enable an audible tone that plays when their Echo device begins and ends streaming audio to the cloud.

(ii) For how long does Alexa listen for subsequent commands after the wake word is spoken when this setting is enabled?

Alexa will end the stream immediately once the user ends the conversation or if Alexa detects silence or speech that isn't intended for Alexa. Customers can find information about the Follow-Up Mode in the Alexa App and on Amazon.com (<https://www.amazon.com/gp/help/customer/display.html?nodeId=202201630>).

Thank you again for your interest in Amazon's privacy and data security practices with respect to our Alexa voice service.

Sincerely,



Brian Huseman
Vice President, Public Policy

¹ Customers can optionally enable an additional feature – Alexa Guard – that enables their Echo device to detect additional specific sounds, such as the sound of smoke alarms or glass breaking, and stream audio of those sounds to the cloud. No audio is sent to the cloud for that feature unless the Echo device detects one of the selected sounds.