



Developer Innovation Toolkit for Voice Enabled Tech, V1

Overview

This document is meant to provide security guidance for developers of voice enabled technologies to embed better security in their products. One of the goals of the EWF Voice Privacy Alliance is to foster innovation and provide open source practical tools to the developer community. These Agile security stories are meant to highlight potential security considerations developers can put into their planning process. In no way are these security stories meant to be prescriptive. The Voice Privacy Alliance hopes that developers will bring this document in-house and adapt it as they see fit based on business considerations and the features and functionality of their innovative products.

We hope the development community will provide feedback and input into future versions of this document. The toolkit is open and available at <http://www.ewf-usa.com/voiceprivacy> or contact us directly at voiceprivacy@ewf-usa.com. To learn more about the Executive Women's Forum (EWF) please go to <http://www.ewf-usa.com>

A special thank you to the following people for taking their valuable time to review and make comments.

Dr. Edward Amoroso, TAG Cyber LLC

Kelly Arnholt, Oracle

Dan Cornell, Denim Group

Dr. Galina Datskovsky, Vaporstream

Kelly Fitzsimmons, Hypervoice Consortium

Maggie Gloeckle

Stacey Gray, Future of Privacy Forum

Rhonda Maclean, Maclean Risk Partners

Yolanda Smith, Pwnie Express

ID	Security Story	Security Guidance	NOTES
33	<p>Customers worry their home wireless network might be insecure and they're worried about someone eavesdropping. Now there's a fear that a hacker could listen in on conversations with their devices.</p>	<p>Consider that voice data could be intercepted from an insecure home network. Even if you are not storing the actual voice data (using speech recognition) customer perceptions that someone else could listen in and record what they say can have a very negative impact on your product and company brand. When protecting voice data in transit you can transmit via TLS – or another acceptably secure protocol – that will protect the data from interception or modification regardless of the configuration of the local home network. A solid overview of using TLS to protect communications can be found in the OWASP TLS Cheat Sheet: https://www.owasp.org/index.php/Transport_Layer_Protection_Cheat_Sheet Overall, it's best to assume the home wireless network is insecure.</p>	<p>Products that need to be manually activated, or for speech recognition use cases where customers will issue commands for relatively benign things such as weather and traffic, may not need this because even if someone was eavesdropping, there may not be high value to the data. Also, while the guidance here is focused on voice enabled technology, existing developer security considerations for applications in customer home networks obviously should be applied.</p>
34	<p>Customers and the press have concerns the device is "always listening" and this makes them uncomfortable, especially if the device is in a home.</p>	<p>Give customers a "mute" capability when a device or application is always listening, even if it's just listening for a wake word (phoneme). And a hardware kill switch is best because hackers could disable a software kill switch. Giving customers this kind of control may help alleviate fears and uncertainty with devices that need to "listen" as critical requirement. You might also consider being transparent with customers how much you buffer when in listen mode. For example, your device might have a 30 second buffer as it listens for a wake word. After 30 seconds the buffer is cleared and the cycle repeats until the customer uses the wake word. Being transparent may alter customer perceptions and make them more comfortable with your product.</p>	<p>Beyond voice specific security concerns, developers also should consider what customer data is collected and stored, including how long the data is stored. With voice activated applications, developers need to consider what resides in a hardware buffer and not just what may be written to disk. Development managers might also want to consult with their legal and regulatory touchpoints because if any stored data, even in a hardware buffer, causes customers significant harm or damage, you might run afoul of regulators such as the FTC (Federal Trade Commission) for unfair deceptive trade practices.</p>

ID	Security Story	Security Guidance	NOTES
35	<p>Customers have read in the press that their voice recordings are stored in the cloud. They ask how are you making sure some bad guy isn't stealing this? Customers feel a voice recording is extremely personal, more so than a social security number or another personal identifier.</p>	<p>Voice data stored on remote servers should be encrypted in transit and at rest, and if not needed long term, data should have a short time-to-live (TTL). Consider that voice data which will be used for speech recognition purposes (e.g., customer asks for the weather, the data is converted to text and the application returns the request) should be discarded as soon as possible to lower the risk and also not take up unnecessary storage space. Developers need to consider that storing a customer's actual voice recording should require encryption at rest, even if it's for a short period of time. There are guidelines already in place which may help developers http://www.biometrics.gov/standards/Registry_v5_2014_08_01.pdf. It's also important to consider server-side services that are used to receive voice recordings - the services should have proper authorization checks to make sure that data can only be retrieved by the intended devices or accounts.</p>	<p>Storing data for short periods of time for speech recognition purposes may not warrant encryption at rest, e.g., a customer asking for the weather forecast. Keep up with voice impersonation attacks and make sure your device can't be used against your customers whose voice data has been captured maliciously</p>
36	<p>Customers fear what they say will live forever. Their concerns might be expressed as, "If you record something I say, can I delete it? What if I said something I'm ashamed of or embarrassed by? Since your device lives in my home or is something I wear, it's possible you'll record something very personal or potentially embarrassing."</p>	<p>If you are storing customer voice data, give them control over their own data. Allow customers the ability to delete their voice data, even if there are UX implications to the product. On the other hand, you should inform the customer how deleting their voice data could negatively impact their user experience or maybe even limit your product's capabilities.</p>	<p>Product managers must consider the impact this kind of functionality will have on the development timeline. In addition, if your product relies on analysis of this voice data there could be financial and even functional impact to the product.</p>

ID	Security Story	Security Guidance	NOTES
37	<p>There are customers who are aware that do not have to associate their voice data with them. Some customers know there are ways to anonymize the data so it's useful to a company but can't be traced back to them.</p>	<p>Anonymize data by default unless there is a business reason to identify and track an individual</p>	<p>Product managers should be aware if their product's financial model relies on tracking individual users and how they are interacting with their devices.</p>
38	<p>Same as #34</p>	<p>In always listening mode, keep voice data cache small and purge often, listening only for the wake word. This reduces the potential attack surface</p>	<p>There is likely a hardware limitation to how big a cache you have to work with.</p>
39	<p>Customers are keenly aware of malware and viruses. Due to recent news, they are also aware the police might take their device and use the data that's on it to prosecute them. What can you do to protect the voice data that resides on your customer's device?</p>	<p>If storing voice data on a customer's local device, use encryption at rest, leveraging existing OS/platform security tools. Like a customer's home network, assume the device is insecure and may have malware. Also, developers will likely need to consider use cases where law enforcement will want to gather voice data from devices.</p>	<p>Encrypting local data at rest can present challenges due to key management issues and may require specialized hardware to implement properly. Product managers also need to be aware of e-discovery implications and should know their company's policy about handing customer data over to law enforcement.</p>
40	<p>Customers are already aware a device could respond to someone else's voice command. There's also the fear, especially with wearables, that voice enabled devices could respond to the local environment and even inadvertently respond to someone</p>	<p>Develop use case scenarios where someone other than an authorized user issues commands to the device. Obviously this does not apply if the device is designed to take input from anyone. This issue is very use case dependent - threat modeling early in the development lifecycle will help identify specific threats that apply to your product/device. Regardless, developers should test all possible inputs as you would normally do.</p>	<p>If your device takes commands from anyone by design, then it's important to inform your customers that's how your product works. Transparency and good documentation can help avoid problems and set proper expectations.</p>

ID	Security Story	Security Guidance	NOTES
	<p>else's unintentional command.</p>		
41	<p>If your device communicates with the customer, how will your customer know it's your device talking and not some interloper like a hacker?</p>	<p>Develop use case scenarios where an attacker might intercept the conversation between a customer and their device, potentially compromising the confidentiality and integrity of the conversation. For example, a voice enabled children's toy is made to say something wholly inappropriate to a child because the communication channel has been compromised. Again refer to #33 above and the use of TLS.</p>	
42	<p>Customers don't want to bring devices into their homes which can communicate with them and potentially their children in inappropriate ways. They also recognize that they don't control the algorithms that go into calculating the appropriate response. This fear is based on a lack of control over how a device or application might respond to them and their family.</p>	<p>If using AI, develop use case scenarios where certain input causes inappropriate system response.</p>	<p>Microsoft Tay is a recent example</p>

ID	Security Story	Security Guidance	NOTES
43	Customers understand that companies share customer data with 3rd parties. Sharing voice data, especially an actual voice recording may cause concern and mistrust.	If you authenticate a customer to a 3rd party service, do so securely and make sure the customer is aware they have crossed over a trust boundary. Again there is a fundamental difference between lower risk voice recognition scenarios where a customer might request a stock quote, you convert that request to text, send it to a third party and the customer gets a response to their request. If you are sending voice data to a 3rd party to analyze, maybe parse and store, then you should consider securing that communication and building in a way for customers to consent.	Product managers should be aware this can be a complex and potentially costly endeavor. Normally, a data mapping exercise should reveal any instances where voice data is transmitted to 3rd parties. In addition, product managers might want to talk with legal and regulatory touchpoints to make sure you are not running afoul of any laws. Further, vendors should also have formal policies in place of how they will handle data that is transmitted to them as 3rd parties. Product managers might consider reaching out to compliance touchpoints to see if they can perform testing/auditing of the adherence to these policies for high-value data.
44	Customers want to know if you are keeping up with current threats and able to update the software/hardware in a timely manner?	Accumulated security debt or vulnerabilities with existing deployed product may require a security sprint. Further, systems should be designed from inception to be updateable – both to accept functional updates as well as security updates.	
45	Customer will ask, “Can my voice-enabled device just respond to me? In my home and with my wearables, I live in a dynamic environment and I wonder if my device could respond to someone else’s command.”	Verify if a single user should have a unique session connection with your product or if multiple sessions are allowed. For example, Hello Barbie may have several people interact with her, but the doll is associated with a child and that child’s authorizing adult. Consider if there are use cases where it would be appropriate to reject sessions not initiated by the authorized user. Amazon Echo is the opposite, designed to respond to anyone who invokes the wake word	There are many ways to authenticate users. Product managers will have to consider the options.

ID	Security Story	Security Guidance	NOTES
47	Customers who worry you are always listening to them will eventually wonder how long you're actually recording and what you're doing with that recording.	In listen mode, document how many seconds the device records prior to hearing the wake word (5 seconds, 30 seconds??) Explain what happens to this recording? Is it stored locally? Is it part of a cache that will be deleted from the hardware? Is it ever sent remotely?	A person's voice is deeply personal and getting in front of concerns like this will help to minimize FUD. In addition, recent consumer press articles seem to like this angle because it captures reader's attention.
48	Customers want to know if you are keeping up with current threats and if you are able to detect new threats on their device. More savvy customers will know most systems have logging capability not only to troubleshoot, but also to detect unusual behavior. Many even send this information voluntarily back to the vendor	Ensure there is balanced overall as well as security event logging for analysis and improvement (better security & privacy understanding)	
49	See #33	If customers authenticate to your device or mobile application, assume weak security on home networks and assume weak passwords (if that's what you're using). Consider leveraging platform security features or establishing an encrypted channel.	
50	After many data breaches, customers know to ask how you are protecting their data from hackers.	If you give customers access to their voice data, consider using OOB authentication or 2-factor authentication. Again assume customers will use weak passwords or that their credentials are already leaked	
51	Customers know that at times they will give an incorrect command or someone else might direct a command to a device that they would like to take back.	If a customer cancels a voice directed action, roll back all changes and purge the voice data locally and/or remotely	

ID	Security Story	Security Guidance	NOTES
52	<p>Customers already have voiced concerns that someone else take over their device and tell it what to do. As wearables become more commonplace, they will worry about being in dynamic open/public places where there are potentially many conversations and voice inputs.</p>	<p>Verify all sessions to ensure no one else can intercept or tamper with a conversation between your product and your customer. It might be helpful in the setup process to guide the customer through the logical steps to secure their device and also understand potential misuse cases. There's a big difference between an app that authenticates for purchases or financial transactions compared to a voice activated TV remote control. A business also needs to consider how much security should come out-of-the-box by default. You can always give customers the option to change security settings but a more secure posture out-of-the-box might be less risky for your product.</p>	<p>Look at both software and hardware weaknesses. For example, devices could send and receive radio waves which may be a way to send a command or take over a device. For example, there was a hack where the headphone connected to your iPhone could activate Siri.</p>
53	<p>Customers understand they're using their voice the way they used to use a keyboard and mouse. There is a use case where customers browse folders with their personal data. They will wonder if they are secure when they do this by voice rather than by keyboard and mouse.</p>	<p>If you give customers the ability to browse directories with voice commands, use the same methods for you would for normal browsing restrictions to ensure they cannot access restricted directories. Make sure that services and devices are authenticating to one another in an appropriate manner.</p>	
54	<p>Customers are well aware of the amount of information that's collected. Collecting unnecessary information is a bad customer experience and erodes trust.</p>	<p>Avoid collecting and storing voice data and metadata unnecessarily. Your application may not have a need to store this kind of data but you might collect it on a one-time basis during setup or to bootstrap critical functionality. It's best to delete this data as soon as its not needed. Similar to #35</p>	<p>Don't collect more data from customers than your business and application require.</p>

ID	Security Story	Security Guidance	NOTES
55	Customers don't want complex choices and UI when trying to configure their security & privacy settings. This is a source of frustration and often leads to less security & privacy.	UI - make sure it's easy for customers to change their security & privacy settings. Consider your application or device and see if there's an opportunity to enforce security & privacy settings by default. Further, clear documentation and infographics are also approaches to educating customers on their security and privacy setting choices.	There will always be a tradeoff between security and privacy settings that protect customers versus providing data valuable to your company
56	Customers know security & privacy are important.	Document user security & privacy settings in your online help, application help system, and in any written documentation that comes with the product. Use all the normal channels your company already uses to communicate with the customer. As a developer it's important that you may need to review for accuracy and be sensitive to the fact that your customers aren't security experts or necessarily technical.	
57	The customer purchased your product and now it's time to do initial setup. The last thing you want to do is give customers a bad user experience as you walk them through the security settings setup. For voice enabled products, there may be specific settings and considerations you will need to communicate clearly to your customer.	Allow customers to set security and privacy settings through automated first time setup routines, much in the way home wireless routers do today using easy-to-understand illustrated instructions. Further, applications that rely on voice for access or authentication should come out-of-the-box with the most stringent security controls enabled. It's always possible to guide the customer through disabling stringent security settings if the customer chooses. This is an especially important step during on-boarding because customers aren't likely to go back in and increase or change security protocols once on boarding is complete.	
58	Customers hate strange error messages. If your device communicates with your customers by speaking, make sure spoken error messages make sense. In today's environment, your	As with all error messages, make sure spoken error messages contain minimal details but is still understandable to customers.	

ID	Security Story	Security Guidance	NOTES
	company could be easily ridiculed for this in social media		
59	Savvy customers are aware that they can be re-directed by hackers and that there are ways to fool them into connecting to a rogue network.	Test use case scenarios where customers can be re-directed outside of your control or domain by malicious users. For example, you may have a minimal networking stack coded in hardware so that customers can connect your device to their home wireless network. You may only use this during the initial setup process. However, developers should think about misuse scenarios where attackers can exploit your minimal network stack or try to fool users to connect to a rogue network, thinking they are connected to you.	This is not just specific to voice enabled technologies
60	See #45	When using voice authentication, limit the number of attempts and use multi-factor authentication for sensitive services. Multiple voice attempts may be due to low audio fidelity, background noise, limitation of the microphone, etc. It might be worthwhile investigating alternate ways to authenticate such as OOB token, facial recognition, fingerprint reader, etc.	Keep in mind there's a big difference between speech commands and speech recognition. Speech as an authentication solution can allow your device to discriminate between users.
61	See #45	For sensitive data & situations, you might consider logging and alerting on a single voice trying to authenticate to multiple resources. If sensitive communications are involved, users should have an easy way to identify and segregate sensitive transactions. However, be aware that not every voice-based transaction will require the most constrained security protocols.	
62	Customer know about 3 rd party risk and know that companies with with other companies to provide a single service. They want to know who is responsible for the security and privacy of their device.	Make sure hardware, mobile application and server side applications use similar/same level of authentication and session management. This is especially important if there are multiple vendors delivering a single user experience. Creating system threat models is a great way to enumerate the various parties responsible for different components of the system and the security assumptions and requirements for each component.	
63	Technically savvy customers are aware that logging off when a device	Make it obvious to customers how to log off, end a voice session or consider setting an automatic logoff or session timeout.	

ID	Security Story	Security Guidance	NOTES
	or system is not needed is a good thing. This may be a carryover from basic security training they received in the workplace		
64	Customers are very accustomed to security updates and many prefer automatic updates for convenience.	Architect a way to update voice enabled hardware and software automatically and use a secured channel. If this has not already been built into the product, it should be a very high priority item in the next sprint.	
65	Customers want to know you are keeping up with current threats. Rogue wireless networks are a known threat and many customers lack the knowledge to protect themselves against this problem.	Ensure your device can verify it's communicating with legitimate devices or networks. It's possible you can be fooled into connecting to a rogue wireless network or your customer can be fooled	
66	Same as #65	Your application should be aware of device connectivity and warn the user accordingly, particularly if sensitive communications are involved. It would be ideal if the application could alert customers that they are connected to an open (and potentially insecure) network and advise them to attach to a more secure network, especially if their personal or sensitive data will be transmitted over the network. Again, using properly-configured TLS to force communications security can help mitigate some of the risks associated with connecting to malicious networks.	
67	See #44	Consider having a special security flag in your bug tracking system for voice related security bugs. This will help you understand voice technology specific security debt	
68	See #44	Ensure voice related incidents are part of your incident response plan. Consider doing a tabletop exercise based on a voice data incident (especially important if voice authentication data is leaked)	

ID	Security Story	Security Guidance	NOTES
69	Most customers are aware of the problems associated with lost or stolen devices.	Consider physical tampering scenarios. Wearables and IoT devices are shrinking in size. Developers must consider physical threats against lost or stolen devices. Look at #63 because an automatic logoff or session timeout may help. If there is valuable and very personal data on the device, consider encryption at rest.	If the device is already GPS enabled, there may be opportunities to leverage it for lost/stolen devices as we've already seen with mobile phones. Also, note that data-at-rest encryption poses key management challenges and the possible need for specialized hardware to properly implement data-at-rest encryption on small devices.
70	Customers are keenly aware that law enforcement sometimes takes data directly from devices (or even metadata). Some are also aware that it's illegal in some US states to record a person without their express permission, and if you do, it's a punishable by law.	Collect all relevant legal and regulatory requirements and enter them into the backlog. No sprint should start until this step has been completed and approved.	
71	See #44	If you interact with other applications, ensure you cannot be misdirected or fooled into communication with a rogue device or application. Authenticate servers prior to communicating sensitive information to them (using TLS is a way to accomplish this)	
72	Customers will want to know how they verify that deleted voice transactions or recordings have actually been purged. They may also wonder what meta-data is kept and how long it's kept.	Be sure that the privacy statement on any app includes a statement on what kind of meta data is kept about each transaction and how long that meta data is kept (example: police don't have to know the content of a conversation in order to make their case--often meta data such as call records or CDRs can be used to reconstruct the conversation)	

ID	Security Story	Security Guidance	NOTES
73	There have been enough data breaches that customers may want to be assured you will let them know if their personal information has been compromised and what you are going to do about it.	Developers + security ops should conduct threat modeling (perhaps via tabletop exercises) to determine the opportunities for a breach, how widespread the breach could be as well as the implications of such a breach; developers should build in controls to identify purge, isolate and notify affected users within a predetermined timeline. See #68 regarding incident response.	