Privacy issues and security risks in Alexa Skills

February 24 2021

Together with colleagues, Christopher Lentzsch has analyzed extra functions for Amazon's voice assistant Alexa. Credit: Ruhr-Universitaet-Bochum

With the voice commands 'Alexa Skills,' users can load numerous extra functions onto their Amazon voice assistant. However, these Skills can

often have security gaps and data protection problems, as a team of researchers from the Horst Görtz Institute for IT Security at Ruhr-Universität Bochum (RUB) and North Carolina State University discovered, together with a former Ph.D. student who started to work for Google during the project. They will present their work at the Network and Distributed System Security Symposium (NDSS) conference on 24 February 2021.

More than 90,000 Skills analyzed

In their study, the research group of Christopher Lentzsch and Dr. Martin Degeling studied first-time the ecosystem of Alexa Skills. These voice commands are developed not only by the U.S. tech company Amazon itself but also by external providers. Users can download them at a store operated by Amazon directly, and in some cases, they are also activated automatically by Amazon.

The researchers obtained and analyzed 90,194 Skills from the stores in seven country platforms. They found significant deficiencies for safe use. "A first problem is that Amazon has partially activated Skills automatically since 2017. Previously, users had to agree to the use of each Skill. Now they hardly have an overview of where the answer Alexa gives them comes from and who programmed it in the first place," explains Dr. Martin Degeling from the RUB Chair of System Security. Unfortunately, it is often unclear which Skill is activated at what time. For example, if you ask Alexa for a compliment, you can get a response from 31 different providers, but it's not immediately clear which one is automatically selected. Data that is needed for the technical implementation of the commands can be unintentionally forwarded to external providers.

Publishing new Skills under a false identity

"Furthermore, we were able to prove that Skills can be published under a false identity. Well-known automotive companies, for example, make voice commands available for their smart systems. Users download these believing that the company itself has provided these Skills. But that is not always the case," says Martin Degeling. Although Amazon checks all Skills offered in a certification process, this so-called Skill squatting, i.e., the adoption of already existing provider names and functions, is often not noticeable.

"In an experiment, we were able to publish Skills in the name of a large company. Valuable information from users can be tapped here," explains the researcher. So if an automotive supplier has not yet developed a Skill for its smart system in the car to turn up or turn down the music in the car, for example, attackers would be able to do so under the supplier's name. "They can exploit users' trust in the well-known name and in Amazon to tap into personal information such as location data or user behavior," Degeling says. Criminals, however, could not directly tap encrypted data or change commands with malicious intent in this process to manipulate the smart car, for example to open the car doors.

Circumventing Amazon's security check

The researchers also identified another security risk: "Our study also showed that the Skills could be changed by the providers afterward," explains Christopher Lentzsch from the RUB Chair of Information and Technology Management. This vulnerability places the security of the previous certification process on the part of Amazon into another perspective. "Attackers could reprogram their voice command after a while to ask for users' credit card data, for example," Lentzsch says. Amazon's testing usually catches such prompts and does not allow them—the trick of changing the program afterward can bypass this control. By trusting the abused provider name and Amazon, numerous users could be fooled by this trick.

Unsufficient data protection declarations

In addition to these security risks, the research team also identified significant lacks in the general data protection declarations for the Skills. For example, only 24.2 percent of the Skills have a so-called Privacy Policy at all, and even fewer in the particularly sensitive areas of "Kids" and "Health and Fitness." "Especially here, there should be strong improvements," Degeling says.

Amazon has confirmed some of the problems to the research team and says it is working on countermeasures.

More information: Christopher Lentzsch, Sheel Jayesh Shah, Benjamin Andow, Martin Degeling, Anupam Das, William Enck: Hey Alexa, is this skill safe?: Taking a closer look at the Alexa Skill ecosystem, Network and Distributed System Security Symposium (NDSS), 2021: www.alexa-skill-analysis.org/#paper

Provided by Ruhr-Universitaet-Bochum