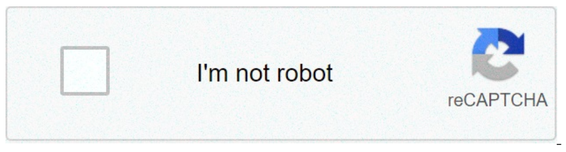

Continue

Contact track and trace

Most people know you can catch a cold from someone who has one. This is why sales of hand sanitizers take off during “cold season” and why people with the sniffles are asked to call in sick or work from home — they hope to minimize the spread of their illness.When a more serious outbreak develops, health authorities may undertake contact tracing or partner notification, a method of identifying and following up with everyone who came into contact with people known to be infected. 1. Origins of Contact Tracing Health authorities have been using contact tracing for decades to help curb the spread of infectious diseases. Companies dedicated to the practice and public health workers travel worldwide, including to Haiti, Peru, and Rwanda, to identify and track many infectious diseases, from tuberculosis and HIV to sexually transmitted diseases. I have been fortunate enough to be part of a Portuguese working group that is actively analyzing and discussing the security and privacy implications of a future contact tracing app in Portugal. In particular, with regard to securing these apps from potential threats. We have seen some countries taking different approaches (centralized, decentralized) and several different technologies being leveraged to build these apps (both native technologies and web technologies like JavaScript).About the authorPedro Fortuna is the Co-Founder and CTO of Jscrambler.Most discussion on contact tracing apps is centered around the issue of privacy, so one of my roles has been to shed light on the underlying security issues that may make it easier for attackers to tamper with contact tracing apps and potentially breach privacy on a massive scale. With several countries all over the world developing and launching contact tracing apps, it is timely to analyse the differences between these apps and to highlight relevant security issues. Let's take a closer look.When it comes to COVID-19 apps, there are 2 types of apps that handle sensitive data - symptom tracking and contact tracing. Contact Tracing is the more controversial of the two because of the fear that people have of being monitored. Non-digital contact tracing has existed for a long time. Once someone has been infected, health officials interview the person to check with whom the person had contact recently. They ask for contact info of people that had contact and interview them as well. This enables tracking contagion and isolating required persons. But the main problem is that this is ad-hoc - people are often interviewed too late in the game and can’t remember everyone who they have been in contact with. Plus, the interviewee may not have their contact info. Another problem is that people are only being isolated after they display symptoms.Digital contact tracing aims to expand upon this and solve some of these limitations. Most contact tracing apps are based on Bluetooth Low Energy (BLE). Each app instance transmits short-range beacons; so, as people move, their devices pick up other people’s beacon. Each beacon has an ephemeral anonymous identifier that is unique. The ephemeral IDs seen are stored locally in the device. When someone is diagnosed, that person receives an official code from health officials, and by entering that code in the app, that person is officially and voluntarily registered with the infection status. Then, a list of its own ephemeral IDs used in the last 14 days will be sent to a central server; all other devices periodically download the list of infected people’s IDs and then locally determine if they have been in contact with them recently. For positive contacts, the application then calculates the risk of infection, taking into consideration the estimated proximity and duration of the contact. If the risk is over a certain threshold, the app triggers a notification recommending that the user self-isolates and contacts the national health authority.Exposed code and reverse-engineeringDue to the fact that they handle very sensitive data, contact tracing apps should offer robust security and be subject to independent security audits. To this end, these apps should follow the general recommendations of the OWASP Mobile Security Testing Guide, specifically the use of Code Signing to reduce the risk of publishing an adulterated version of the application, and also Certificate Pinning, to reduce the risk of Man-in-the-Middle (MITM) attacks. Additional security threats that should be taken into account can be found here.Of particular note here is the risk posed by exposed code. As outlined by OWASP, this is where an attacker can either directly modify the code, change the contents of memory dynamically, change or replace the system APIs that the application uses or modify the application’s data and resources. Subsequently, this can provide the attacker with a direct method of subverting the intended use of the software for personal or monetary gain. The impact from code modification can be wide ranging in nature, depending upon the nature of the modification itself. But it is especially relevant in decentralized Contact Tracing apps where the data and sensitive algorithms all stay inside the local environment.OWASP also highlights the risk of attackers reverse-engineering the source code as a gateway for more advanced attacks. In this way, an attacker may exploit reverse engineering to achieve any of the following aims - to reveal information about backend servers, to reveal cryptographic constants and ciphers, to steal intellectual property, to perform attacks against back end systems or to gain intelligence needed to perform subsequent code modification.Framing this into contact tracing apps, both these security risks (exposed code & reverse engineering) must be addressed in order to block this security gap. This issue is especially important in JavaScript-based apps - where, by default, the source code is not compiled and is therefore completely exposed, greatly facilitating tampering and reverse-engineering. Currently, there are dozens of different contact tracing apps either under development or released. We can see that more than a few of these are being developed using JavaScript-based frameworks such as React Native and Ionic - namely, the official Israeli app Hamagen, Germany-based ito, and Canada-based COVID Shield (which are all based on the React Native framework) and Switzerland’s WeTrace which is built with the Ionic framework.Accordingly, it is critical that these apps protect their JavaScript source code from the attacks described previously, with a special focus on protection against data exfiltration attacks that can be achieved by manipulating code or by memory inspection. To this end, OWASP advises that the mobile app must be able to detect at runtime that code has been added or changed, and that the app must be able to react appropriately at runtime to a code integrity violation and must protect itself from memory tampering or scraping. They go on to state that in order to prevent effective reverse engineering, a code protection tool must be used.With the teams behind these projects facing the giant challenge of delivering these highly complex apps in record time, it’s essential that they consider these security threats still during the development stage. This need for source code protection in mobile apps mustn’t be ignored with stakes as high as these. To ensure that their apps are protected against attacks to the integrity of the source code, these teams must look for resilient JavaScript code protection. We’ve made our coronavirus coverage free for all readers. To get all of HBR’s content delivered to your inbox, sign up for the Daily Alert newsletter. Efforts to slow the spread of Covid-19 in the U.S have been stymied by the lack of an effective national surveillance system that can track the emergence of suspected and confirmed new cases in real-time. This is in part because important patient data is trapped in siloed electronic health record systems that don’t communicate well with each other; a new case of suspected Covid-19 in one health system could be invisible to another in the same community, blinding the systems and public health officials to Covid-19 hotspots that may be developing right in front of them. If we could see such hotspots as they appear we could direct resources where they’re needed most to treat new patients and contain transmission. Leveraging an existing platform that already accesses the electronic health records of 200,000 physicians and health care providers across more than 400 hospitals nationally, our firm has built an automated, real-time surveillance app that integrates with existing electronic health records (EHRs) from different companies. This solution could provide early warning capability, forecast surges, and help providers plan coordinated responses. The EHR-agnostic app will be offered for free this year to the dozens of health systems that use our clinical platform beginning with Geisinger Health System in July. Atrium Health, Community Health Network, Vidant Health, AdventHealth and other health systems have confirmed that they plan to rapidly follow. We initially assumed that the CDC or other government agencies already had this type of capability. But we discovered that existing “syndromic surveillance,” which tracks emerging infectious disease and other public-health issues, is not automated and does not consistently provide the complete information needed to support robust containment and mitigation strategies. As Kaiser Health News reported, clinicians and public health officials sometimes must print out data from EHRs and manually fill in and fax reporting forms. Some CDC syndromic surveillance forms can take up to 30 minutes to complete, and the lag between a patient receiving a positive Covid-19 test and the reporting of that data can be as long as seven days. CDC syndromic surveillance systems, in fact, may have provided false reassurance by reporting an artificially low rate of infection through the ILI-NET surveillance program in mid-April 2020 in California, Florida, and Michigan, even though the virus was running rampant in those states at the time. This potential blind spot may have occurred because surveillance primarily tracked emergency department patients who, fearing infection, had started to avoid ED visits. Without an effective national or state system for disease surveillance and monitoring, the U.S. response to Covid-19 and future pandemics will continue to face handicaps and have critical knowledge gaps. For instance, today we continue to lack clear, real-time, nationwide data detailing: Where and in which settings patients are presenting with Covid-19-like symptoms; The local hospitals that are likely to admit positive cases after they are detected; Risk- and severity-adjusted information that allows provider systems to predict the supplies they will need to care for their specific patient populations; Which patients are likely to become sicker over time; How physicians are caring for positive patients, including therapies and prescription medications used in treatment; Whether care provided to patients with Covid-19 is consistent with the latest scientific evidence; and, Where the disease is surging in real time. Our application can’t yet do all of this, but it’s a start and it demonstrates how a surveillance app integrated onto disparate existing EHRs could. To the best of our knowledge ours is the first such app to show this capability. More than a dozen existing commercial products that leverage similar functionality and EHR access could be repurposed to do this. The ability of our app to access cloud-based patient records in real-time across EHR systems depends on a little-known federal law called PAMA (The Protecting Access to Medicare Act of 2014). PAMA requires that most U.S. physicians use digital tools that enable the review of patient information in EHRs as the patient is being treated. Our app integrates with a decision-support tool that is used by more than 200,000 physicians and other clinicians as they order medical imaging procedures across nearly 35 health systems. When imaging is ordered, the app — with appropriate permission — can access the patient record. The app uses natural language processing and machine learning to scan clinicians’ free-text patient records and orders for terms including “trouble breathing” and “loss of sense of taste” amongst many other Covid-19-related signs, symptoms, and other indications of infection. In tests, we have been able to rapidly identify patients who are presenting with signs and symptoms associated with Covid-19 syndrome. By subsequently associating these flagged patients with their Covid-19 test results we believe that the app can be trained to be highly sensitive and specific, identifying infected patients with a relatively low rate of false positives and negatives. Early testing has shown our symptom false-positive rate to be approximately 5%, which we expect to improve as the software continues to learn. We are now working to expand the app’s access beyond the narrow pool of patients receiving imaging orders to all patients receiving Covid-19 tests, and, from there, we will work to expand to a broader group of patients that have not yet been tested. Data will be collected when patients are evaluated by telehealth technology, in physicians’ offices, in urgent care and in emergency departments. This information is accessible through major EHRs and includes patients’ symptoms, signs, medications, laboratory tests, and other data. The version that will survey all patients receiving Covid-19 tests, with or without an imaging study, will be rolled out in July 2020, before a potential second wave forecasted for fall or winter of 2020. Ultimately, when the app’s surveillance is expanded to the broader population, it will predict patients’ risk status and, with the ability to detect upticks in suspected cases from a 7-day moving average, the need for hospital beds. This forecasting ability will also allow providers to schedule elective surgeries and procedures to avoid Covid-patient surges. Because central coordination of Covid-19 surveillance data is largely occurring at a state level, we are in discussions with state and federal officials to send data feeds from health systems to state health departments to automate what is mostly an antiquated and manual process. Patient privacy will continue to be carefully protected as manual processes become automated. The passage of the PAMA law in 2014 has serendipitously enabled the rapid development of a potentially effective, real-time Covid-19 surveillance system. The required EHR platforms have already been installed in virtually every physician’s office and, as we have shown, straightforward modification of existing apps running on those systems can be readily rolled out. This approach has enormous potential for tracking the current pandemic, detecting subsequent waves of Covid-19, and modernizing population-level surveillance of emerging infectious disease. If our free content helps you to contend with these challenges, please consider subscribing to HBR. A subscription purchase is the best way to support the creation of these resources.