CYBER
THREAT
ANALYSIS

Recorded Future®

By Insikt Group®

April 16, 2024

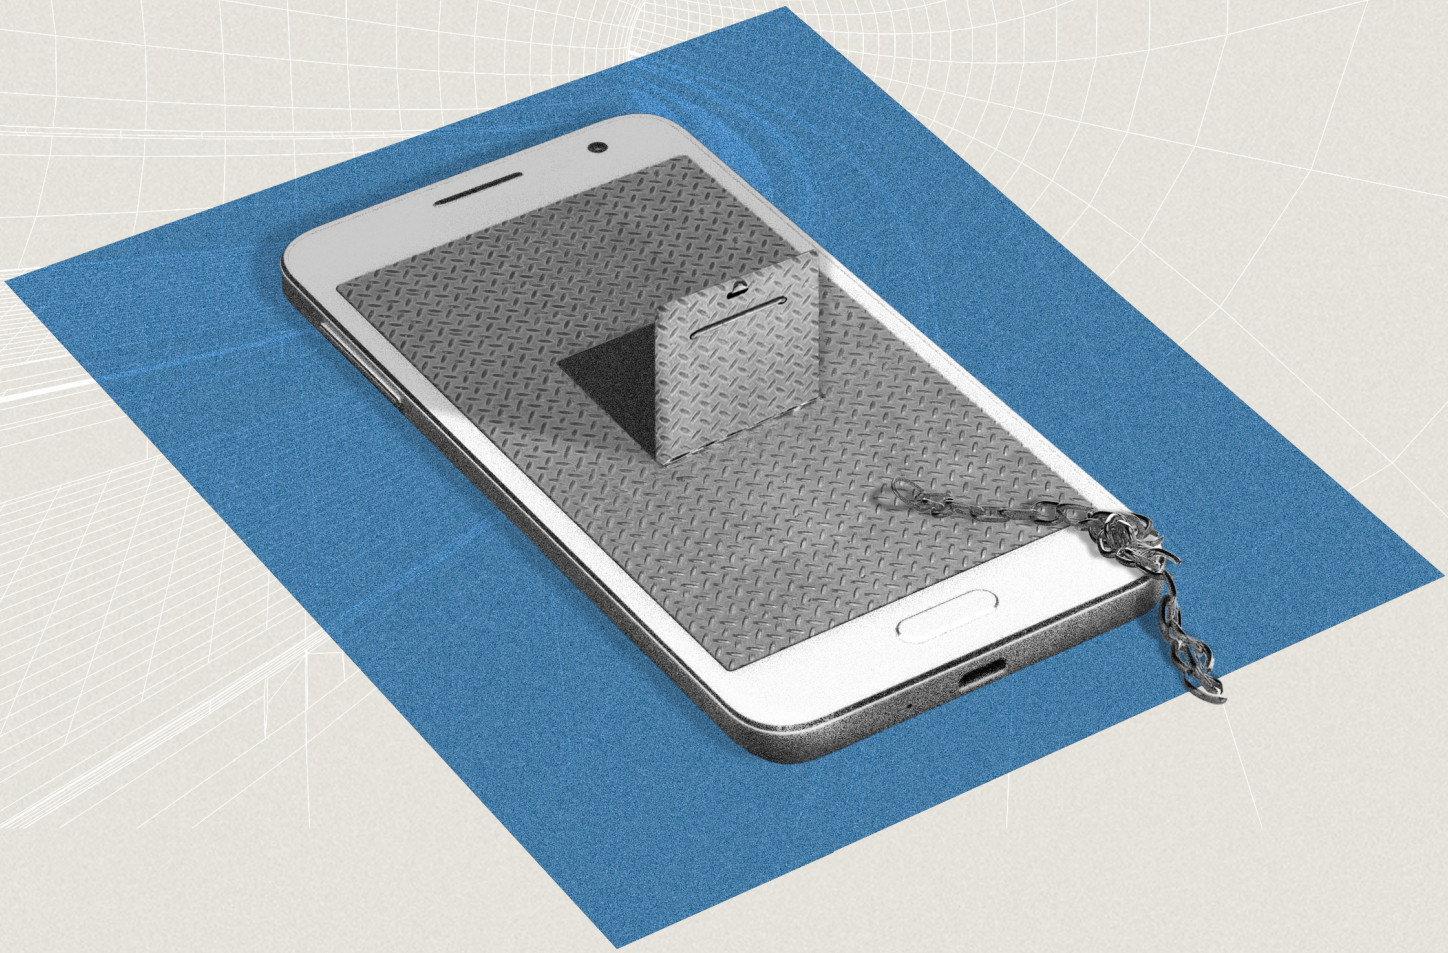

“Mobile NotPetya”: Spyware Zero-Click Exploit Development Increases Threat of Wormable Mobile Malware

Executive Summary

In 2017, the NotPetya ransomware cost organizations worldwide billions of dollars due to its ability to spread quickly and at scale. The likelihood of a "mobile NotPetya" event, characterized by a self-propagating mobile malware deliberately or accidentally spread through zero-click exploits, is increasing rapidly. This escalation is mainly driven by the persistent development of zero-click exploits by spyware companies. In 2023 alone, more zero-click vulnerabilities were disclosed than in the prior four years combined (six vs. three). The likelihood of this rate of disclosure continuing into the next few years is high based on an ongoing increase of mobile users, increasing integration of mobile devices into corporate security (such as with app-based multi-factor authentication), and the increasing reliance on mobile devices for other corporate activities. Additionally, the motivations or conditions for releasing a mobile NotPetya into the wild have never been greater due to worldwide geopolitical flashpoints and the sophistication of cybercriminal groups. These conditions allow for several scenarios involving wormable mobile malware, including deliberate deployment by nation-states as part of a hybrid assault and accidental deployment by criminal groups as part of the implementation of exploits into existing mobile malware campaigns. The impact of a mobile NotPetya is also likely to be severe, with the potential for hundreds of thousands of devices being compromised within the first few days of deployment.

Ten years ago, wormable mobile malware variants could spread rapidly but could also be easily mitigated due to the more simplistic nature of telecommunications at the time. By contrast, current mobile device messaging across both iOS and Android platforms involves numerous complicated processes, protocols, integrations, and file formats that make protection from zero-click exploits much harder for both telecommunications companies and mobile device manufacturers. Some aspects of those platforms, such as shared messaging across Apple devices, introduce an additional risk beyond mobile devices. There are potential “emergency brakes”, such as Lockdown Mode for iOS, that companies could use to stop a mobile NotPetya, but there is as yet no general public reporting or systematic testing of the efficacy of these measures.

Key Findings

- All of the conditions necessary for a mobile NotPetya event are already in place. These include the ongoing development of zero-click exploits, the ability for mobile malware to abuse contact lists for further spread, a “monoculture” landscape of two primary mobile device operating systems, a lack of clear mitigations from telecommunications and mobile device companies, and a heightened risk of nation-state actors deploying zero-click exploits against targets due to geopolitical conflict.
- The two most noteworthy historical wormable pieces of mobile malware are Commwarrior and Cabir, both of which relied on Nokia phones to spread. The first was highly narrow in its targeting, and the second could be blocked at the email gateway level, so neither achieved the spread that could have affected their full populations of hundreds of millions of users.

- Three separate zero-click cyber threat campaigns (such as Operation Triangulation) involving the exploitation of six zero-day vulnerabilities have been disclosed within two quarters in 2023, representing a jump from any previous news about zero-click exploitation over the previous four years. We are concerned about this trendline because the more such vulnerabilities are available to attackers, the higher the likelihood that they will be abused, deliberately or not, to infect victims at scale.
- Multiple mobile device manufacturers have released what are effectively “safe modes” for text messaging to thwart zero-click exploits (one example is Apple’s Lockdown Mode). These safe modes work largely by reducing the multimedia functionality of messages, which has been a core component of previous exploits. However, such steps are unlikely to be very effective since they are either unlikely to appeal to a large group of mobile users (especially younger or technically unsavvy users) or are limited in what they can prevent.
- Research into how the spread of malware compares to the spread of disease has resulted in multiple mathematical models that help predict the effects of a mobile NotPetya event. One highly relevant model, for example, predicts hundreds of thousands of mobile device infections within the first few days of a campaign relying on zero-click exploits and contact list abuse for spreading.
- We believe that there are at least two places where a telecommunications provider or mobile device manufacturer could stop a wormable mobile malware in its tracks: by filtering messages based on header information that is consistent across malicious messages or by filtering messages based on the geographic location of clusters of victim devices. However, neither of these measures has been tested at scale, and we suspect that their implementation would only occur when infection rates have already risen past the point of easy containment, making them less effective mitigations.

Threat/Technical Analysis

What Is "Mobile NotPetya"?

In 2017, the NotPetya malware infected hundreds of thousands of devices and became the most costly cyberattack ever conducted, with some companies suffering losses of over $500 million. The malware largely targeted Windows systems that were vulnerable to the ETERNALBLUE exploit.

As of this writing, no mobile malware event has been as severe or as widespread as the NotPetya campaign. However, the chances of such an event — a "mobile NotPetya" in which wormable mobile malware spreads indiscriminately with severe consequences — have increased, in some aspects exponentially, over the last ten years.

NotPetya's severity was due to its combination of two threats: the ETERNALBLUE exploit itself, which abused the common networking protocol SMB; and the open-source credential dumping tool Mimikatz. Together, ETERNALBLUE and Mimikatz allowed NotPetya to propagate across networks at scale and without an operator's direction.

Similarly, we anticipate that a mobile NotPetya event, in which mobile malware spreads in a self-propagating fashion to smartphone users, requires that the malware be able to autonomously *infect* targets easily and then *identify new* targets easily. The most likely vectors for each, respectively, are an exploit for a zero-click vulnerability in a smartphone operating system and a function to abuse victims' contact lists for further spread. One-click vulnerabilities (such as those requiring that victims click a link in an SMS message) would be much less effective since there would be a higher detection rate.

As was true for NotPetya, we assume that the most likely threat actor to be responsible for a mobile NotPetya event would be a nation-state-resourced group and that the most likely reason for such a group's deployment would be destructive. Given the latter assumption, we consider it extremely unlikely that a hacktivist group would be responsible. However, a cybercriminal group (such as one already invested in spreading mobile malware) could also be responsible.

In the "Mobile NotPetya" Recipe, All Ingredients Are Available

For a high-severity wormable mobile malware scenario to occur, the following items are required:

- Development of zero-click exploits for one or more zero-day vulnerabilities affecting one of the two major mobile operating systems, Android and iOS
- The ability for mobile malware to spread autonomously (such as via abuse of contact lists, Bluetooth, and so on)
- A landscape in which zero-day exploitation of a single vendor has outsize effects; in other words, many users are likely to be affected based on market saturation of particular operating systems.

- Lack of effective "emergency brake" mitigations from telecommunications providers or mobile device manufacturers
- Motivation for threat actors to dedicate resources to identify vulnerabilities, develop exploits, and develop mobile malware

In several cases, these items have already been in play for ten years or more. At the very least, the conditions for all of these to coincide as a result of a motivated threat actor have been increasingly likely due to several developments in the past decade:

- **Vulnerability research:** Zero-day, zero-click vulnerabilities affecting Android and iOS continue to be developed or paid for by spyware vendors like the NSO Group. Organizations like Citizen Lab have been able to identify and disclose these vulnerabilities and associated exploits, but their work likely only covers a fraction of the total landscape of mobile device vulnerability exploitation.
- **Contact list abuse:** Mobile malware has had the ability to spread autonomously for many years. As early as 2012, the mobile malware FireLeaker was observed stealing victims' contacts to pre-position for spam campaigns. Bluetooth has also been a vector for autonomous spread in mobile malware campaigns from twenty years ago (see Background section below).
- **Disproportionate effects:** Between them, Android and iOS account for a user base of about five billion users, or about 93% of an estimated 5.4 billion individual smartphone users worldwide. While there are many versions of each OS available (or not updated) among this user base, these numbers still represent a victim landscape with "monoculture" vulnerabilities in the face of single threats.
- **Unclear mitigations:** While we suspect that the mobile device manufacturing and telecommunications industries have thought about this as a problem scenario, our research uncovered no public policy or statements from these industries about how they are equipped to stop or mitigate a mobile NotPetya event, and we were unsuccessful in soliciting feedback from individual companies.
- **Heightened motivations:** There are now multiple areas around the globe where geopolitical tensions have exploded into ongoing kinetic conflict or are likely to do so in the next few years. The cybercriminal ecosystem is also now dominated by tightly run or high-volume operations servicing large-scale, costly threat campaigns. Between these two, there are many potential motivations for unleashing a wormable mobile malware.

Out of all of the factors above, observers of the mobile threat landscape should be most attuned to the motivation aspect. We do not consider it to be hyperbole or fear-mongering to imagine a scenario, for example, in which a Chinese invasion of Taiwan includes wormable mobile malware as an initial salvo in disrupting communications for Taiwanese and allied military units. There are also no inherent barriers to a cybercriminal group gaining access to a spyware developer's zero-day exploit and incorporating it into a mobile malware campaign without fully understanding or caring about its severity. Moreover, given the secrecy of spyware developers, a breach of their exploit capabilities might not be known until much later than a similar breach of a more public entity.

Background: Previous Wormable Mobile Malware Reliant on Bluetooth to Spread

Wormable mobile malware is not a theoretical concept. The first smartphone virus ever to be deployed was Cabir, a worm that affected Symbian-based Nokia devices in 2004. Cabir abused Bluetooth to distribute malicious Software Installation Script (SIS) files, which were installer files used by the Symbian OS. Cabir infections were seen in twenty countries and led to what was described as a "mass outbreak" in at least one location, Finland, after a sporting event in Helsinki. In its retrospective on the malware, Kaspersky said that it was "very easy" to become infected with Cabir due to the virus's abuse of Bluetooth, although neither Kaspersky nor other sources have released an estimate of the final victim count for the malware.

A similar malware variant, Commwarrior, followed shortly after. Like Cabir, it sought to distribute malicious SIS files. In 2005, it had "spread to eight countries" and was particularly dangerous because it was an early abuser of Multimedia Messaging Service (MMS) messages; unlike Cabir, Commwarrior could read and send further malicious messages to individuals in an infected device's contact list.

Fortunately for users, while the attack surfaces for Cabir and Commwarrior were trivial for criminals to exploit, they were equally simple to mitigate. Cabir could be stopped completely if a user turned off Bluetooth, and its abuse of emerging smartphone email functionality could be identified and stopped at the cell carrier level. Commwarrior was even less of a threat because it was reportedly "fussy" about its targets (for example, only affecting the Nokia Series 60).

In the wake of malware like Cabir and Commwarrior, and almost certainly in response to smartphone usage increasing, mobile operating system and device manufacturers hardened devices against such simplistic exploits. As a result, in the last decade, mobile malware has been far more likely to spread via social engineering than via vulnerabilities. However, there are already indicators that threat actors are seeking new vectors for the spread of their malware. As recently as May 2023, Trend Micro identified a malicious operation from "Lemon Group" that had pre-installed the "Guerilla" malware on nearly nine million Android devices via mobile supply-chain compromise; similarly, the BADBOX fraud campaign has relied on the compromise of "physical off-brand Android devices".

Exploits Abuse Increasingly Complex Mobile Messaging Functionality

It is likely that when cybersecurity-conscious people hear about exploitation of smartphone vulnerabilities, they suspect a small set of spyware developers who have become prominent in the past few decades for this type of work. This is an appropriate connection to make: the most well-known spyware companies, like Israel's NSO Group, have been responsible for the ongoing development of smartphone vulnerability exploits, with NSO Group typically (but not exclusively) focusing on iOS exploits. This business can be hugely profitable, especially if a company has flexible limits on which countries or organizations it is willing to sell to. We are almost certain that dozens of government agencies worldwide, regardless of geography or type of government, continue to sponsor this industry's

attempts to exploit smartphone technology. We also note that NSO Group is only one of many spyware developers whose activities are typically covert and difficult to analyze.

A Rising Trend of Zero-Click Exploits

A list of zero-click mobile device vulnerabilities and their associated exploits over the past several years shows that threats of this kind are not only not going away but have consistently been able to bypass defense mechanisms set up to thwart them. It is the "zero-click" aspect of these vulnerabilities that is most problematic, removing any inherent defenses that still exist against social engineering attacks (not everyone will click on a link). Additionally, in some cases, victims of these exploits have been targeted indiscriminately, countering any argument that zero-click vulnerabilities would only be exploited for singular, high-value compromises due to their high cost.

- In May 2019, WhatsApp discovered that Pegasus spyware was being delivered to victims, including a "UK-based human rights lawyer", via a zero-click vulnerability (CVE-2019-3568) in WhatsApp 2.19.99. Attackers could deliver malware to victims by calling a WhatsApp number and delivering a series of malicious RTP Control Protocol (RTCP) packets.
- In November 2019, Motherboard obtained and shared a copy of the FTI Consulting report detailing how Amazon CEO Jeff Bezos's iPhone may have been compromised via malicious WhatsApp messages from Saudi Crown Prince Mohammed bin Salman. The description of the hack matches that of one using a zero-click exploit, but FTI Consulting could not identify details of any malware or vulnerabilities involved.
- In February 2020, Android disclosed and patched a zero-click Bluetooth vulnerability affecting Android 9, identified as CVE-2020-0022. There was no news at the time or since of exploitation in the wild.
- In May 2020, Google disclosed a zero-click vulnerability (CVE-2020-8899) that affected all Samsung smartphones sold since 2014, up to Android 10. Samsung patched the vulnerability in the same month. Setting the tone for future mobile exploits, this zero-click vulnerability resided in an image processing component, specifically "how the Android OS flavor running on Samsung devices handle[d] the custom Qmage image format (.qmg)". However, the researcher who identified the vulnerability noted that executing a zero-click exploit for this would require 50-300 MMS messages to discover the vulnerable area, which would take about 100 minutes and likely alert a victim.
- In December 2020, Citizen Lab published a report on an iOS zero-click exploit it called "KISMET", which affected the messaging component of iOS 13.5. In particular, KISMET abused iCloud partitions and a "background process that appear[ed] to be associated with iMessage and FaceTime". The exploit had been used to infect the devices of dozens of Al Jazeera personnel in the summer of 2020 with NSO Group's Pegasus spyware.
- In September 2021, Citizen Lab published a report on an iOS zero-click exploit it called "FORCEDENTRY", which, again, affected the messaging component of an Apple device (iOS 14.7). In this case, the malicious operation used files with fraudulent `.gif` extensions to hide malicious Adobe PDF and PSD files. It worked by abusing a vulnerability, CVE-2021-30860, in

Apple's image rendering library. Citizen Lab discovered the exploit being used to infect the device of an unnamed Saudi activist.
- In April 2022, Citizen Lab published a report on a spyware campaign targeting a wide range of individuals associated with the Catalan independence movement. This campaign exploited an iOS zero-click vulnerability the group called "HOMAGE", which involved the launching of a WebKit instance via iMessage. The campaign also included the use of the KISMET exploit (see above).

We started writing this report in February 2023 based on an assumption that the zero-click vulnerability landscape was likely to worsen over the subsequent year, even with the release of security measures like BlastDoor for iOS 14. A set of vulnerability disclosures in 2023 — two of them high-profile — confirmed this assumption:

- In April 2023, Apple disclosed and patched a use-after-free vulnerability (CVE-2023-28205) affecting multiple iOS versions up to iOS 15. The company said that the vulnerability "may have been actively exploited" but provided no further details. However, since Amnesty International was involved in the vulnerability's discovery, external observers speculated that it had been exploited to compromise individuals involved in human rights advocacy. As has been true for other zero-click iOS vulnerabilities, it was part of Apple's WebKit browser engine.
- In June 2023, Kaspersky disclosed that it had discovered exploitation of a zero-click vulnerability in iOS 16.1, CVE-2022-46690, as part of a campaign that the company dubbed "Operation Triangulation". Like KISMET and FORCEDENTRY, the exploit abused the WebKit module and the "iCloud" name for spoofed domains, respectively. On the same day of Kaspersky's disclosure, the Russian government accused Apple of colluding with the US government to compromise "several thousand" Russian iPhone users (neither the Russian government nor any other researchers have since released data to back up this accusation) In December 2023, Kaspersky released further details of the campaign, finding that it relied on a sophisticated multi-phase exploit that abused six separate vulnerabilities, including four that were zero-days at the time of infection. The most notable aspect of this exploit chain was a vulnerability that could be abused to write data to "unknown hardware registers of the chip unused by the firmware"; Kaspersky confessed that it had "no idea" how attackers would have identified or known how to exploit this vulnerability.
- In September 2023, Citizen Lab broke the news that a threat campaign had used a zero-click, zero-day exploit to bypass BlastDoor's defenses (hence its name, "BLASTPASS") and distribute Pegasus spyware. Uncharacteristically, Citizen Lab released few technical details of the exploit, and so other researchers' attempts to explore the vulnerability have been merely guesswork.

The examples above disproportionately mention iPhones, but spyware developers have their eyes on any and all operating systems. Also, in September 2023, *Haaretz Magazine* published a report about an Israeli firm that had developed a tool called Sherlock that could infect users of all operating systems via malicious advertisements. These are not technically zero-click exploits (a victim would still need to visit a site with ads), but the report noted a worrying trend in the spyware community:

> Until now, different companies have specialized in breaching different devices. Candiru focused on PCs, NSO could hack iPhones, and its competitors specialized in Androids. But with this system, as the documents show, every device could effectively be breached.

If we only chart those zero-click smartphone vulnerabilities that were known to have been exploited in the wild, the timeline since 2019 is as follows:

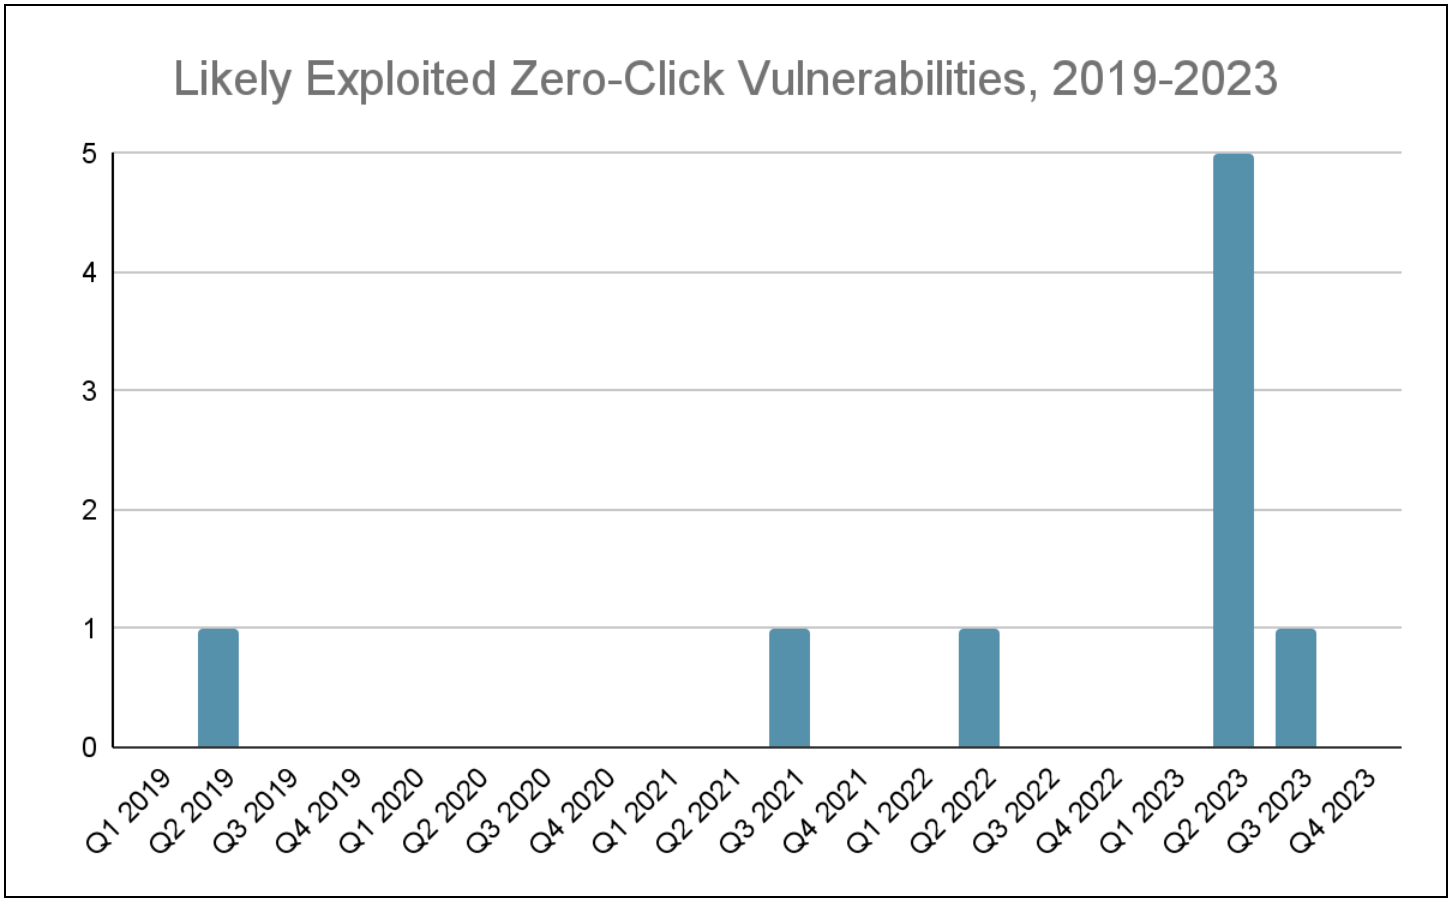

Figure 1: *Zero-click vulnerability exploitation began to proliferate at an alarming rate in 2023 compared to the previous four years*

The news of three separate zero-click campaigns involving the exploitation of six zero-day vulnerabilities within the span of two quarters represents a huge jump from any previous news about zero-click exploitation over the past four years, let alone the entire time frame of mobile exploitation. We are very concerned about this trendline because the more such vulnerabilities are available to attackers, the higher the likelihood that they will be abused, deliberately or not, to infect victims at scale.

Current Defenses Are Inadequate Due to MMS Complexity and Users' Preferences

Part of the difficulty for smartphone manufacturers' defense against exploit developers is the complexity demanded of modern mobile messaging. The term "text message" for mobile

communications undersells the amount of functionality that users expect from their smartphones. Over the last few decades, mobile messages have developed from mere text strings (part of the original Short Message Service [SMS] protocol) to clusters of media components that require an array of protocols to be encapsulated, transmitted, received, and reconstituted. As with many other forms of software, this explosion in complexity has allowed a concomitant rise in difficult-to-detect vulnerabilities that threat actors can exploit to gain access to devices.

Multiple mobile device manufacturers have taken steps to make text messaging more resistant to zero-click exploits. For example, in September 2022, Apple released its Lockdown Mode, which heavily restricts a multitude of applications (such as those for messaging, web browsing, and photos) to reduce a target's attack surface. (Apple later claimed that Lockdown Mode would have protected against the BLASTPASS exploit that Citizen Lab disclosed.) Then in February 2023, Samsung unveiled its Samsung Message Guard, an "isolated virtual space" that examines files to see if they hide malicious code.

Both of these cases are good examples of mobile device manufacturers taking steps to protect their users, and both have almost certainly made their users safer. However, like any other security measure, they should not be seen as absolutely effective. In particular, Lockdown Mode requires that users opt out of an entire ecosystem of interconnected multimedia functionality that Apple almost certainly depends on to attract customers, especially younger customers.

In many cases, and again fortunately for users, mobile operating system source code is kept secret, and even when it is leaked, it does not always include its own legend to help outside observers understand how different processes and referents fit together (although this information is also sometimes exposed). However, since companies like NSO Group have clientele that can afford multimillion-dollar exploits, they have been able to continue allocating funds and expertise toward compromising these "black-box" systems.

In December 2020, Ian Beer, a researcher for Google's Project Zero, published an article about how he was able to find a "zero-click radio proximity exploit" for iOS. His observations throughout his lengthy breakdown of the discovery are worth reading, but two stand out in particular for a discussion of defenses:

- As demonstrated by Beer's discussion of it, the kernelcache for a modern iPhone represents a huge number of functions that Apple has to confirm are safe, on their own and integrated with other applications, with every release of a new version. This situation applies to Android operating systems as well. Protecting this space can be a cat-and-mouse game since people have been publishing reverse engineering tools (such as this one from 2016) for Apple's kernelcache for many years.
- Disclosing and fixing a vulnerability often invites an increased risk of exploits: "It's well known to all vulnerability researchers that the easiest way to find a new vulnerability is to look very closely at the code near a vulnerability which was recently fixed. They are rarely isolated incidents and usually indicate a lack of testing or understanding across an entire area".

Malware and Disease: Epidemiological Models Show How Mobile Malware Could Spread

Based on a review of relevant mathematical models, a wormable mobile malware event could feasibly compromise hundreds of thousands of mobile users, for either iOS or Android, within a few days. Statistically speaking, the chance that such an event would affect both major operating systems at the same time is, of course, even lower than the chance of such an event happening at all.

To reach this conclusion, we needed to reference a field that has already generated useful predictions for real-world events involving "wormable" harms. In this regard, articulating a precise estimate of malware dispersal benefits greatly from the field of epidemiology, which similarly needs to be able to forecast how a negative condition spreads within a population.

In particular, it is the "zero-click" aspect of recent mobile vulnerabilities that has made comparisons with epidemiology appropriate. In a typical mobile malware campaign (such as for FluBot), victims must be tricked into clicking links for malicious websites or downloads. However, in a wormable mobile malware campaign, victims would be immediately exposed and compromised if their device featured a certain vulnerability, which is much more similar to a scenario in which a new disease infects a population that has no innate immunity.

Fortunately, the cybersecurity research community has explored the comparisons between the spread of malware and the spread of disease for over a decade. In 2009, for example, the US National Academy of Sciences (NAS) released a report on "WiFi Networks and Malware Epidemiology", which found that "tens of thousands of [vulnerable] routers" could be infected by a malware strain "in as little as two weeks, with the majority of the infections occurring in the first 24–48 h". The "tens of thousands" number came from looking only at the router density in Manhattan (which the researchers estimated at over 35,000), which means the number in a globally relevant scenario would likely be much larger.

The idea of certain devices being more or less susceptible to infection is a theme that reappears in subsequent papers comparing cybersecurity and epidemiology. The NAS report differentiated routers based on how easily they could be infected (for example, routers with default passwords vs. simple passwords vs. highly complex passwords). While this theme is important to consider for a mobile NotPetya event, it is not as relevant since such an event would be severe precisely because devices would have no built-in mitigations against zero-click, zero-day vulnerability exploitation.

One of the most relevant pieces of research for a mobile NotPetya event was published in December 2020 by the Institute of Electrical and Electronics Engineers (IEEE) with the title "Stochastic Modeling of IoT Botnet Spread: A Short Survey on Mobile Malware Spread Modeling". This research was prompted by the spread of the Mirai botnet on Internet-of-Things (IoT) devices. Since "IoT botnets aggressively scan for IP addresses to randomly find their victims", the researchers saw Mirai as analogous to an infectious disease, which also spreads to victims aggressively and randomly. The thirteen mathematical

models discussed by the researchers are essential for coming to an understanding of how severe a mobile NotPetya scenario could be.

In many cases, the epidemic model of "Susceptible-Infection-Recovery" (S-I-R) has been used to describe the movement of malware through a population of devices. This model can be helpful since it takes into account how malware might be removed from a system in the course of a major cyber threat campaign. However, as the IEEE researchers point out, "IoT devices are less likely to gain immunity and are continuously vulnerable to new infections … think of this as the flu season during which everyone is susceptible to the flu virus, irrespective of having a flu vaccine". A zero-click mobile malware event fits exactly into this paradigm of no available immunity and difficulties in recovery.

One of the mathematical models referenced in the report (in subsection B) applies to the spread of malware via Bluetooth (BT). This formula was used to "accurately estimate the distribution curve of BT worms in large cities like Los Angeles". We used this model for the global population of iPhone users in the Washington, DC, area alone (~2.8 million), with a rate of two days for infected devices to be removed from the network. Based on this model, **within three days of just one initially infected device, hundreds of thousands of devices would be infected.**

- We note that if, as we have hypothesized, a wormable mobile malware campaign would take advantage of users' contact lists to spread, there is a high chance that at least some malware distribution would result in "closed loops" in which infected devices could only reach out to already infected devices.
- The total population of iPhone users is approximately 1.5 billion, and that of Android users is about 3.3 billion. If these victim populations were affected by mobile malware at the rates forecast by the model above, the results would be devastating. Depending on how these operating system manufacturers delivered a patch, the number of susceptible devices could still be very high; for example, in a 2016 survey of smartphone users in the US, only about 30% of respondents said that they used iPhone's auto-update feature to stay updated. That said, we consider it very likely that in a mobile NotPetya scenario, companies like Apple or Google would be highly incentivized to deliver non-voluntary patches.

Mitigations: How "Emergency Brakes" Could Shut Down a Mobile NotPetya

Given the increasing number of zero-click exploits that have been disclosed over the past few quarters, the large attack surface within the MMS ecosystem, and the unlikely scenario of a majority of mobile users adopting features like Lockdown Mode, the landscape for a mobile NotPetya event can be only minimally improved by the actions of individual or corporate mobile users. The best that can be done at such a level is to ensure that devices are in some form of restricted mode, such as Lockdown mode for iOS. This is especially important for mobile devices associated with critical business permission or operations or for mobile devices that are provided to employees.

A more effective method for shutting down a mobile NotPetya requires potentially disruptive measures on the part of telecommunications providers, mobile device manufacturers, or both.

There is precedent for telecommunications providers having tools at hand to stop the spread of wormable mobile malware. As previously mentioned, the Cabir mobile worm could be stopped at the cellular carrier level insofar as it relied on malicious emails to spread. However, a mobile malware variant with the ability to spread worm-like via a zero-click exploit would be unlike any prior mobile malware in terms of mitigation strategies.

At present, we believe that there are at least two places where a telecommunications provider or mobile device manufacturer could stop a wormable mobile malware in its tracks: by filtering messages based on header information that is consistent across malicious messages or by filtering messages based on the geographic location of clusters of victim devices. However, neither of these measures has been tested at scale, and we suspect that their implementation would occur when infection rates have already risen past the point of easy containment.

Here it is useful to briefly explain the structure and delivery of an MMS message. Unlike SMS messages, MMS messages need to be able to encapsulate and transfer information regarding a variety of media types, including photos, videos, audio, `.gif` files, device-specific stickers, and so on. Also, unlike SMS messages, which were based on radio wave protocols to send nothing but text strings, MMS messages rely on TCP/IP networks for faster and more efficient delivery of these additional media types. Even when MMS messages are using radio-wave protocols like General Packet Radio Service, they are still doing so in order to send IP packets. As with other communications within TCP/IP, MMS messages are encapsulated into an encoded message plus a header that helps with filtering and routing the message to its final recipient.

Here is the first place where a telecommunications provider or mobile device manufacturer might be able to stop a mobile NotPetya event. There are many variables in the header for an MMS message (such as sender, recipient, timestamp, message class, message structure, and so on), but one that might remain unchanged across a malicious campaign is a message's content type and *location*. The

emphasis on location is deliberate: in many cases, when a user sends an MMS message, their device is not itself compressing media information into packets for full transfer; rather, their device includes referents to externally hosted media that the receiving device can then access and incorporate as part of the reassembly of the message inside a messaging application. This means that if a zero-click exploit campaign relies on a single malicious referent to infect other devices (such as a fake PDF or `.gif` file), a telecommunications provider or mobile device manufacturer should be able to access either the content of a message itself or its media referents to create deny-list filters for mobile devices within their networks.

The second vector for shutting down a mobile NotPetya would be much less refined and much more disruptive. Telecommunications providers can shut off MMS connectivity for specific geographies within their purview if they receive news of a wormable mobile malware scenario developing regionally. A recent example of a telecommunications shutdown occurring involuntarily was in November 2023, when, as a result of the kinetic conflict between Israel and Hamas, Gazan telecommunications companies Paltel and Jawwal announced an impending blackout for their user base. More generally, governments around the world have already enforced telecommunications shutdowns as part of efforts to "quell mass protests, forestall election losses, reinforce military coups, or cut off conflict areas from the outside world". While these are often illegitimate uses of this measure, they demonstrate that it is a viable option against a mobile NotPetya.

A final emergency brake against a mobile NotPetya could simply be operator error. The severity of the Wannacry ransomware campaign, for example, was famously mitigated by researchers sinkholing a domain that the malware used for communications. A mobile NotPetya might similarly rely on an easily stoppable piece of infrastructure, or it might fizzle out due to the malware being badly designed and therefore failing to spread past a certain number of devices. Trusting this type of mitigation is, of course, wishful thinking.

Outlook

A mobile NotPetya event is not destined to happen. There are many variables associated with security research, mobile device configuration, the zero-click exploit market, and geopolitics that could reduce or eliminate the conditions necessary for a mobile NotPetya event to occur in the near future, if at all. However, at present, we are living in a world in which all of the conditions for mobile NotPetya are not just active but are worsening each year, and therefore we must take them seriously.

From our perspective, one of the greatest dangers associated with this threat landscape is the lack of concerned and emphatic voices in the cybersecurity space talking about the increase in zero-click exploits and explaining how the telecommunications or smartphone industries are equipped to minimize the risk of a wormable campaign relying on them. In other words, the problem now is a lack of attention to the *operational* hazards of the spyware industry. The *moral* hazards of zero-click exploits in spyware campaigns have already been well articulated by groups like Citizen Lab:

> It is now well established that NSO Group, Candiru, other companies like them, as well as their various ownership groups, have utterly failed to put in place even the most basic safeguards against abuse of their spyware.

There are two final quotes that are instructive for thinking through the implications of a currently unlikely — but increasingly likely — and high-severity threat event like a mobile NotPetya.

The first comes from a post on the Graham Cluley blog about the ZCryptor ransomware variant in May 2016, which attempted to spread by copying itself onto any removable drive on a victim system:

> A new ransomware variant exhibits worm-like behavior … ZCryptor might be a harbinger of threats to come.

A year later, NotPetya became the costliest malware attack — not just the costliest encryption attack — of all time.

The second quote comes from a retrospective from Kaspersky (ironically the target of recent smartphone zero-click exploitation) ten years after the appearance of the first smartphone virus, Cabir:

> A few months before Cabir, Alex Gostev, Kaspersky Lab chief malware researcher, was asked by a journalist, why there are no smartphone viruses. Gostev responded that in one year's time there will be some. It turned out, he was right.

Recorded Future®

About Insikt Group®

Recorded Future's Insikt Group, the company's threat research division, comprises analysts and security researchers with deep government, law enforcement, military, and intelligence agency experience. Their mission is to produce intelligence that reduces risk for clients, enables tangible outcomes, and prevents business disruption.

About Recorded Future®

Recorded Future is the world's largest threat intelligence company. Recorded Future's Intelligence Cloud provides end-to-end intelligence across adversaries, infrastructure, and targets. Indexing the internet across the open web, dark web, and technical sources, Recorded Future provides real-time visibility into an expanding attack surface and threat landscape, empowering clients to act with speed and confidence to reduce risk and securely drive business forward. Headquartered in Boston with offices and employees around the world, Recorded Future works with over 1,700 businesses and government organizations across more than 75 countries to provide real-time, unbiased, and actionable intelligence.

Learn more at recordedfuture.com