How we protect users from 0-day attacks

blog.google/threat-analysis-group/how-we-protect-users-0-day-attacks/

Maddie Stone July 14, 2021

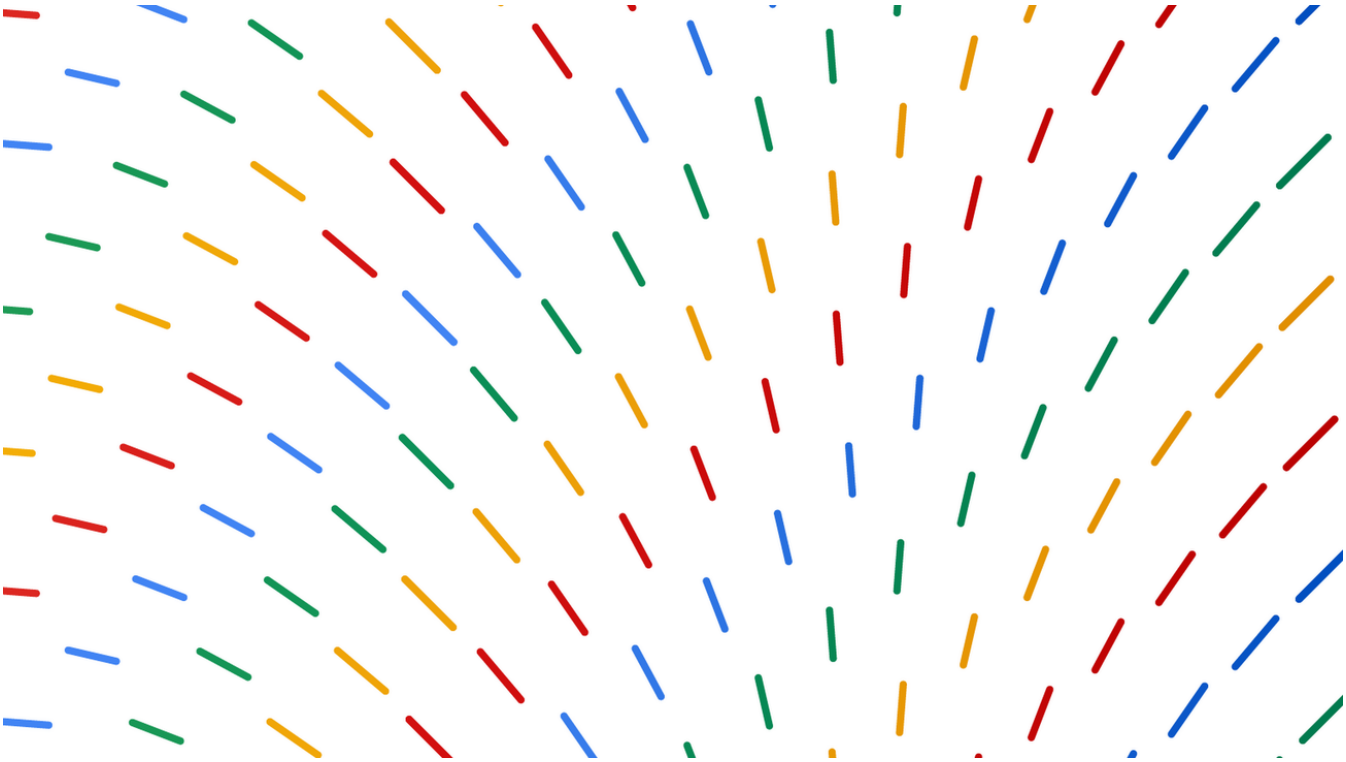

Zero-day vulnerabilities are unknown software flaws. Until they're identified and fixed, they can be exploited by attackers. Google's Threat Analysis Group (TAG) actively works to detect hacking attempts and influence operations to protect users from digital attacks, this includes hunting for these types of vulnerabilities because they can be particularly dangerous when exploited and have a high rate of success.

In this blog, we're sharing details about four in-the-wild 0-day campaigns targeting four separate vulnerabilities we've discovered so far this year:

- CVE-2021-21166 and CVE-2021-30551 in Chrome,
- CVE-2021-33742 in Internet Explorer, and
- CVE-2021-1879 in WebKit (Safari).

The four exploits were used as a part of three different campaigns. As is our policy, after discovering these 0-days, we quickly reported to the vendor and patches were released to users to protect them from these attacks. We assess three of these exploits were developed by the same commercial surveillance company that sold these capabilities to two different government-backed actors. Google has also published root cause analyses (RCAs) on each of the 0-days.

In addition to the technical details, we'll also provide our take on the large uptick of in-the-wild 0-day attacks the industry is seeing this year. Halfway into 2021, there have been 33 0-day exploits used in attacks that have been publicly disclosed this year — 11 more than the total number from 2020. While there is an increase in the number of 0-day exploits being used, we believe greater detection and disclosure efforts are also contributing to the upward trend.

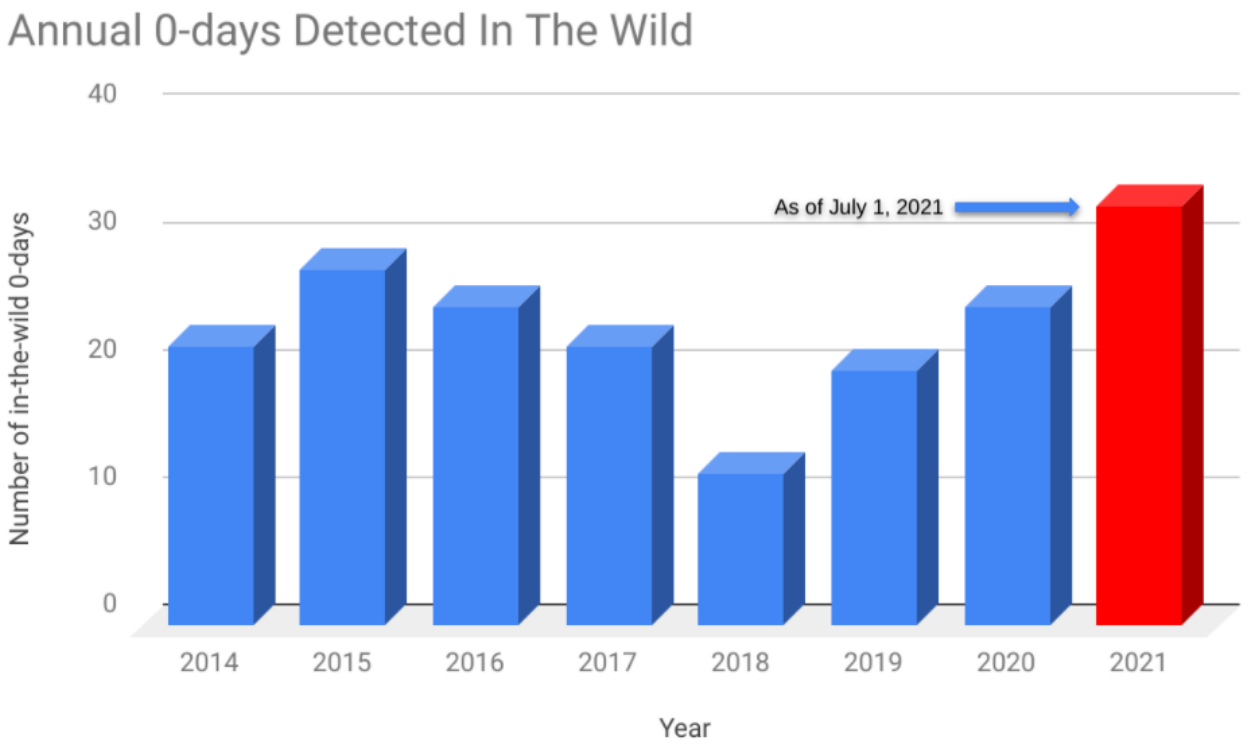

Chrome: CVE-2021-21166 and CVE-2021-30551

Over the past several months, we have discovered two Chrome renderer remote code execution 0-day exploits, CVE-2021-21166 and CVE-2021-30551, which we believe to be used by the same actor. CVE-2021-21166 was discovered in February 2021 while running Chrome 88.0.4323.182 and CVE-2021-30551 was discovered in June 2021 while running Chrome 91.0.4472.77.

Both of these 0-days were delivered as one-time links sent by email to the targets, all of whom we believe were in Armenia. The links led to attacker-controlled domains that mimicked legitimate websites related to the targeted users. When a target clicked the link, they were redirected to a webpage that would fingerprint their device, collect system information about the client and generate ECDH keys to encrypt the exploits, and then send this data back to the exploit server. The information collected from the fingerprinting phase included screen resolution, timezone, languages, browser plugins, and available MIME types. This information was collected by the attackers to decide whether or not an exploit should be delivered to the target. Using appropriate configurations, we were able to recover two 0-day exploits (CVE-2021-21166 & CVE-2021-30551), which were targeting the latest versions of Chrome on Windows at the time of delivery.

After the renderer is compromised, an intermediary stage is executed to gather more information about the infected device including OS build version, CPU, firmware and BIOS information. This is likely collected in an attempt to detect virtual machines and deliver a tailored sandbox escape to the target. In our environment, we did not receive any payloads past this stage.

While analyzing CVE-2021-21166 we realized the vulnerability was also in code shared with WebKit and therefore Safari was also vulnerable. Apple fixed the issue as CVE-2021-1844. We do not have any evidence that this vulnerability was used to target Safari users.

Related IOCs

- lragir[.]org
- armradio[.]org

- asbares[.]com
- armtimes[.]net
- armlur[.]org
- armenpress[.]org
- hraparak[.]org
- armtimes[.]org
- hetq[.]org

Internet Explorer: CVE-2021-33742

Despite Microsoft announcing the retirement of Internet Explorer 11, planned for June 2022, attackers continue to develop creative ways to load malicious content inside Internet Explorer engines to exploit vulnerabilities. For example, earlier this year, North Korean attackers distributed MHT files embedding an exploit for CVE-2021-26411. These files are automatically opened in Internet Explorer when they are double clicked by the user.

In April 2021, TAG discovered a campaign targeting Armenian users with malicious Office documents that loaded web content within Internet Explorer. This happened by either embedding a remote ActiveX object using a Shell.Explorer.1 OLE object or by spawning an Internet Explorer process via VBA macros to navigate to a web page. At the time, we were unable to recover the next stage payload, but successfully recovered the exploit after an early June campaign from the same actors. After a fingerprinting phase, similar to the one used with the Chrome exploit above, users were served an Internet Explorer 0-day. This vulnerability was assigned CVE-2021-33742 and fixed by Microsoft in June 2021.

The exploit loaded an intermediary stage similar to the one used in the Chrome exploits. We did not recover additional payloads in our environment.

During our investigation we discovered several documents uploaded to VirusTotal.

Based on our analysis, we assess that the Chrome and Internet Explorer exploits described here were developed and sold by the same vendor providing surveillance capabilities to customers around the world. On July 15, 2021 Citizen Lab published a report tying the activity to spyware vendor Candiru.

Related IOCs

Examples of related Office documents uploaded to VirusTotal:

- https://www.virustotal.com/gui/file/656d19186795280a068fcb97e7ef821b55ad3d620771d42ed98d22ee3c635e67/detection
- https://www.virustotal.com/gui/file/851bf4ab807fc9b29c9f6468c8c89a82b8f94e40474c6669f105bce91f278fdb/detection

Unique URLs serving CVE-2021-33742 Internet Explorer exploit:

- http://lioiamcount[.]com/IsnoMLgankYg6/EjlYIy7cdFZFeyFqE4IURS1
- http://db-control-uplink[.]com/eFe1J00hISDe9Zw/gzHvIOlHpIXB
- http://kidone[.]xyz/VvE0yYArmvhyTl/GzV

Word documents with the following classid:

{EAB22AC3-30C1-11CF-A7EB-0000C05BAE0B}

Related infrastructure:

- workaj[.]com
- wordzmncount[.]com

WebKit (Safari): CVE-2021-1879

Not all attacks require chaining multiple 0-day exploits to be successful. A recent example is CVE-2021-1879 that was discovered by TAG on March 19, 2021, and used by a likely Russian government-backed actor. (NOTE: This exploit is not connected to the other three we've discussed above.)

In this campaign, attackers used LinkedIn Messaging to target government officials from western European countries by sending them malicious links. If the target visited the link from an iOS device, they would be redirected to an attacker-controlled domain that served the next stage payloads. The campaign targeting iOS devices coincided with campaigns from the same actor targeting users on Windows devices to deliver Cobalt Strike, one of which was previously described by Volexity.

After several validation checks to ensure the device being exploited was a real device, the final payload would be served to exploit CVE-2021-1879. This exploit would turn off Same-Origin-Policy protections in order to collect authentication cookies from several popular websites, including Google, Microsoft, LinkedIn, Facebook and Yahoo and send them via WebSocket to an attacker-controlled IP. The victim would need to have a session open on these websites from Safari for cookies to be successfully exfiltrated. There was no sandbox escape or implant delivered via this exploit. The exploit targeted iOS versions 12.4 through 13.7. This type of attack, described by Amy Burnett in Forget the Sandbox Escape: Abusing Browsers from Code Execution, are mitigated in browsers with Site Isolation enabled such as Chrome or Firefox.

Related IOCs

- supportcdn.web[.]app
- vegmobile[.]com

  111.90.146[.]198

Why So Many 0-days?

There is not a one-to-one relationship between the number of 0-days being used in-the-wild and the number of 0-days being detected and disclosed as in-the-wild. The attackers behind 0-day exploits generally want their 0-days to stay hidden and unknown because that's how they're most useful.

Based on this, there are multiple factors that could be contributing to the uptick in the number of 0-days that are disclosed as in-the-wild:

Increase in detection & disclosure

This year, Apple began annotating vulnerabilities in their security bulletins to include notes if there is reason to believe that a vulnerability may be exploited in-the-wild and Google added these annotations to their Android bulletins. When vendors don't include these annotations, the only way the public can learn of the in-the-wild exploitation is if the researcher or group who knows of the exploitation publishes the information themselves.

In addition to beginning to disclose when 0-days are believed to be exploited in-the-wild, it wouldn't be surprising if there are more 0-day detection efforts, and successes, occurring as a result. It's also possible that more people are focusing on discovering 0-days in-the-wild and/or reporting the 0-days that they found in the wild.

Increased Utilization

There is also the possibility that attackers are using more 0-day exploits. There are a few reasons why this is likely:

- The increase and maturation of security technologies and features mean that the same capability requires more 0-day vulnerabilities for the functional chains. For example, as the Android application sandbox has been further locked down by limiting what syscalls an application can call, an additional 0-day is necessary to escape the sandbox.
- The growth of mobile platforms has resulted in an increase in the number of products that actors want capabilities for.
- There are more commercial vendors selling access to 0-days than in the early 2010s.
- Maturing of security postures increases the need for attackers to use 0-day exploits rather than other less sophisticated means, such as convincing people to install malware. Due to advancements in security, these actors now more often have to use 0-day exploits to accomplish their goals.

Conclusion

Over the last decade, we believe there has been an increase in attackers using 0-day exploits. Attackers needing more 0-day exploits to maintain their capabilities is a good thing — and it reflects increased cost to the attackers from security measures that close known vulnerabilities. However, the increasing demand for these capabilities and the ecosystem that supplies them is more of a challenge. 0-day capabilities used to be only the tools of select nation states who had the technical expertise to find 0-day vulnerabilities, develop them into exploits, and then strategically operationalize their use. In the mid-to-late 2010s, more private companies have joined the marketplace selling these 0-day capabilities. No longer do groups need to have the technical expertise, now they just need resources. Three of the four 0-days that TAG has discovered in 2021 fall into this category: developed by commercial providers and sold to and used by government-backed actors.

Meanwhile, improvements in detection and a growing culture of disclosure likely contribute to the significant uptick in 0-days detected in 2021 compared to 2020, but reflect more positive trends. Those of us working on protecting users from 0-day attacks have long suspected that overall, the industry detects only a small percentage of the 0-days actually being used. Increasing our detection of 0-day exploits is a good thing — it allows us to get those vulnerabilities fixed and protect users, and gives us a fuller picture of the exploitation that is actually happening so we can make more informed decisions on how to prevent and fight it.

We'd be remiss if we did not acknowledge the quick response and patching of these vulnerabilities by the Apple, Google, and Microsoft teams.

POSTED IN:

Threat Analysis Group