High-severity vulnerability in vBulletin is being actively exploited

reader comments

29 with 25 posters participating

Share this story

Attackers are mass-exploiting an anonymously disclosed vulnerability that makes it possible to take control of servers running vBulletin, one of the Internet's most popular applications for website comments. Sites running the app should take comments offline until administrators install a patch that vBulletin developers released late Wednesday morning.

The vulnerability was disclosed through an 18-line exploit that was published on Monday by an unidentified person. The exploit allows unauthenticated attackers to remotely execute malicious code on just about any vBulletin server running versions 5.0.0 up to 5.5.4. The vulnerability is so severe and easy to exploit that some critics have described it as a back door.

"Essentially, any attack exploits a super simple command injection," Ryan Seguin, a research engineer at Tenable, told Ars. "An attacker sends the payload, vBulletin then runs the command, and it responds back to the attacker with whatever they asked for. If an attacker issues a shell command as part of the injection, vBulletin will run Linux commands on its

host with whatever user permissions vBulletins' system-level user account has access to." Seguin has more in this technical analysis of the vulnerability.

According to researcher Troy Mursch of the Bad Packets security intelligence service, attackers are using botnets to actively exploit vulnerable servers. Some of the Web requests they send look like this:

"widgetConfig[code]=echo shell_exec('sed -i \'s/eval(\$code);/if (isset(\$_REQUEST[\"epass\"]) \&\& \$_REQUEST[\"epass\"] == \"2dmfrb28nu3c6s9j\") { eval(\$code); }/g\' includes/vb5/frontend/controller/bbcode.php && echo -n exploited | md5sum'); exit;"

Some of the infected computers carrying out the attacks have been spotted in the past using the EternalBlue exploit, developed by and later stolen from the National Security Agency, to compromise computers that have yet to install a patch Microsoft released in early 2017.

Some vBulletin users took to the software's official support pages on Wednesday to report they had been hacked. "I received an email today from my hosting provider stating that 'malicious code was detected on your website and a huge number of email spam messages originating from it,'" one user wrote here (free account required). Another user reported having an entire MySQL database deleted.

vBulletin is among the most widely used website commenting systems and is probably used on tens of thousands—possibly hundreds of thousands—of sites. Fortunately, version 5x makes up less than 7% of active installations, according to W3techs, a site that surveys the software used across the Internet. Still, Internet searches like this one suggest that 10,000 or more sites may be running vulnerable versions.

Exploit available for years

According to Chaouki Bekrar, founder and CEO of the Zerodium exploit broker, the vulnerability has been privately circulating for years.

> *The recent vBulletin pre-auth RCE 0day disclosed by a researcher on full-disclosure looks like a bugdoor, a perfect candidate for @PwnieAwards 2020. Easy to spot and exploit.*
>
> *Many researchers were selling this exploit for years. @Zerodium customers were aware of it since 3 years*
>
> *— Chaouki Bekrar (@cBekrar) September 25, 2019*

"Many researchers were selling this exploit for years," he wrote on Twitter. "Zerodium customers were aware of it since 3 years."

The availability of a working exploit is aggravated by another publicly posted script that uses the Shodan search site to find vulnerable servers. Attackers can use it to generate a list of vBulletin sites that are susceptible and then use the exploit to take them over.

The vulnerability exists in default installations of the affected versions. According to Tenable's publicly posted analysis, "an unauthenticated attacker can send a specially crafted HTTP POST request to a vulnerable vBulletin host and execute commands. These commands would be executed with the permissions of the user account that the vBulletin service is utilizing. Depending on the service user's permissions, this could allow complete control of a host."

As advised earlier, the vulnerability is so severe that vulnerable vBulletin users should take their forums offline until they have installed a patch developers published on Wednesday morning. The commenting system for Defcon.org, a

site that’s regularly probed for easy-to-hack vulnerabilities, was non-operational at the time this post went live. This Internet cache shows the site used a vulnerable version as recently as Tuesday. The shutdown likely means that site admins heeded the advice to take comments offline until the vulnerability is fixed.

Before a patch was available, people reported that they were able to successfully mitigate the vulnerability by following the instructions here. Now that a patch is available, affected vBulletin users should install it at once.

https://arstechnica.com/?p=1574761

Apple takes flak for disputing iOS security bombshell dropped by Google

reader comments

178 with 110 posters participating

Share this story

Apple is taking flak for disputing some minor details of last week’s bombshell report that, for at least two years, customers’ iOS devices were vulnerable to a string of zeroday exploits, at least some of which were actively exploited to install malware that stole location data, passwords,

encryption keys, and a wealth of other highly sensitive data.

Google's Project Zero said the attacks were waged indiscriminately from a small collection of websites that "received thousands of visitors per week." One of the five exploit chains Project Zero researchers analyzed showed they "were likely written contemporaneously with their supported iOS versions." The researcher's conclusion: "This group had a capability against a fully patched iPhone for at least two years."

Earlier this week, researchers at security firm Volexity reported finding 11 websites serving the interests of Uyghur Muslims that the researchers believed were tied to the attacks Project Zero identified. Volexity's post was based in part on a report by TechCrunch citing unnamed people familiar with the attacks who said they were the work of nation—likely China—designed to target the Uyghur community in the country's Xinjiang state.

Breaking the silence

For a week, Apple said nothing about any of the reports. Then on Friday, it issued a statement that critics are characterizing as tone-deaf for its lack of sensitivity to human rights and an overfocus on minor points. Apple officials wrote:

> *Last week, Google published a blog about vulnerabilities that Apple fixed for iOS users in February. We've heard from customers who were concerned by some of the claims, and we want to make sure all of our customers have the facts.*
>
> *First, the sophisticated attack was narrowly focused, not a broad-based exploit of iPhones "en masse" as described. The attack affected fewer than a dozen websites that focus on content related to the Uighur community. Regardless of the scale of the attack, we take the safety and security of all*

> *users extremely seriously.*
>
> *Google's post, issued six months after iOS patches were released, creates the false impression of "mass exploitation" to "monitor the private activities of entire populations in real time," stoking fear among all iPhone users that their devices had been compromised. This was never the case.*
>
> *Second, all evidence indicates that these website attacks were only operational for a brief period, roughly two months, not "two years" as Google implies. We fixed the vulnerabilities in question in February – working extremely quickly to resolve the issue just 10 days after we learned about it. When Google approached us, we were already in the process of fixing the exploited bugs.*
>
> *Security is a never-ending journey and our customers can be confident we are working for them. iOS security is unmatched because we take end-to-end responsibility for the security of our hardware and software. Our product security teams around the world are constantly iterating to introduce new protections and patch vulnerabilities as soon as they're found. We will never stop our tireless work to keep our users safe.*

One of the things most deserving of criticism was the lack of sensitivity the statement showed for the Uyghur population, which over the past decade or longer has faced hacking campaigns, internment camps, and other forms of persecution at the hands of the Chinese government. Rather than condemning an egregious campaign perpetrated on a vulnerable population of iOS users, Apple seemed to be using the hacking spree to assure mainstream users that they weren't targeted. Conspicuously missing from the statement was any mention of China.

Nicholas Weaver, a researcher at UC Berkeley's International Computer Science Institute, summed up much of this criticism

by tweeting: “The thing that bugs me most about Apple these days is that they are all-in on the Chinese market and, as such, refuse to say something like ‘A government intent on ethnic cleansing of a minority population conducted a mass hacking attack on our users.’”

> *The thing that bugs me most about Apple these days is that they are all-in on the Chinese market and, as such, refuse to say something like “A government intent on ethnic cleansing of a minority population conducted a mass hacking attack on our users.” https://t.co/ACMhtpN53H*
>
> *— Nicholas Weaver (@ncweaver) September 6, 2019*

The statement also seemed to use the fact that “fewer than a dozen” sites were involved in the campaign as another mitigating factor. Project Zero was clear all along that the number of sites was “small” and they had only a few thousand of visitors each month. More importantly, the size of the campaign had everything to do with decisions made by the attackers and little or nothing to do with the security of iPhones.

Two months or two years?

One of the few factual assertions Apple provided in the statement is that the websites were probably operational for only about two months. A careful parsing of the Project Zero report shows researchers never stated how long the sites were actively and indiscriminately exploiting iPhone users. Rather, the report said, an examination of the five attack chains made up of 14 separate exploits suggested that they gave the hackers the ability to infect fully up-to-date iPhones for at least two years.

These points prompted satiric tweets similar to this one from Juan Andrés Guerrero-Saade, a researcher at Alphabet-owned

security firm Chronicle: "'It didn't happen the way they said it happened, but it happened, but it wasn't that bad, and it's just Uyghurs so you shouldn't care anyways. No advice to give here. Just move along.'"

> *Wow @apple…*
>
> *'It didn't happen the way they said it happened, but it happened, but it wasn't that bad, and it's just Uyghurs so you shouldn't care anyways. No advice to give here. Just move along.'*
>
> *— J. A. Guerrero-Saade (@juanandres_gs) September 6, 2019*

Satire aside, Apple seems to be saying that evidence suggests that the sites that Google found indiscriminately exploiting the iOS vulnerabilities were operational for only two months. Additionally, as reported by ZDNet, a researcher from security firm RiskIQ claims to have uncovered evidence that the websites didn't attack iOS users indiscriminately, but rather only visitors from certain countries and communities.

If either of those points are true then it's worth taking note, since virtually all media reports (including the one from Ars) have said sites indiscriminately did so for at least two years. Apple had an opportunity to clarify this point and say precisely what it knows about active use of the five iPhone exploit chains Project Zero found. But Friday's statement said nothing about any of this, and Apple representatives didn't respond to a request to comment for this post. A Google spokesman said he didn't know precisely how long the small collection of websites identified in the report were operational. He said he'd try to find out, but didn't respond further.

In a statement, Google officials wrote: "Project Zero posts technical research that is designed to advance the understanding of security vulnerabilities, which leads to

better defensive strategies. We stand by our in-depth research which was written to focus on the technical aspects of these vulnerabilities. We will continue to work with Apple and other leading companies to help keep people safe online."

A missed opportunity

Former NSA hacker and founder of the firm Rendition Infosec Jake Williams told Ars that ultimately, the time the exploit sites were active is immaterial. "I don't know that these other 22 months matter," he explained. "It feels like their statement is more of a straw man to deflect away from the human rights abuses."

Also missing from Apple's statement is any response to the blistering criticism the Project Zero report made of Apple's development process, which the report alleges missed vulnerabilities that in many cases should have been easy to catch with standard quality-assurance processes.

"I'll investigate what I assess to be the root causes of the vulnerabilities and discuss some insights we can gain into Apple's software development lifecycle," Project Zero researcher Ian Beer wrote in an overview of last week's report. "The root causes I highlight here are not novel and are often overlooked: we'll see cases of code which seems to have never worked, code that likely skipped QA or likely had little testing or review before being shipped to users."

Another key criticism is that Apple's statement has the potential to alienate Project Zero, which according to a Google spokesman has to date privately reported more than 200 vulnerabilities to Apple. It's easy to imagine that it wasn't easy for Apple to read last week's deep-dive report publicly documenting what is easily the worst iOS security event in its 12-year history. But publicly challenging a key ally on such minor details with no new evidence does not create the best optics for Apple.

Apple had an opportunity to apologize to those who were hurt, thank the researchers who uncovered systemic flaws that caused the failure, and explain how it planned to do better in the future. It didn't do any of those things. Now, the company has distanced itself from the security community when it needs it most.

https://arstechnica.com/?p=1563931

Exploit for wormable BlueKeep Windows bug released into the wild

Aurich Lawson

reader comments

33 with 23 posters participating, including story author

Share this story

For months, security practitioners have worried about the public release of attack code exploiting BlueKeep, the critical vulnerability in older versions of Microsoft Windows that's "wormable," meaning it can spread from computer to computer the way the WannaCry worm did two years ago. On Friday, that dreaded day arrived when the Metasploit framework—an open source tool used by white hat and black hat hackers alike—released just such an exploit into the wild.

The module, which was published as a work in progress on Github, doesn’t yet have the polish and reliability of the EternalBlue exploit that was developed by the NSA and later used in WannaCry. For instance, if the people using the new module specify the wrong version of Windows they want to attack, they’ll likely wind up with a blue-screen crash. Getting the exploit to work on server machines also requires a change to default settings in the form of a registry modification that turns on audio sharing.

By contrast, the wormable EternalBlue exploit—which a still-unidentified group calling itself the Shadow Brokers released into the wild in April 2017—worked seamlessly against a wide range of Windows versions in their default settings. A month after the leak, EternalBlue was folded into the Wannacry ransomware worm that shut down computers worldwide. A month later, another EternalBlue-driven attack called NotPetya created still more worldwide destruction. The latest flaw, which is indexed as CVE-2019-0708 but is better known by the name BlueKeep, resides in earlier versions of the Remote Desktop Services, which help provide a graphical interface for connecting to Windows computers over the Internet. It affects Windows 2003 and XP, Vista 7, Server 2008 R2, and Server 2008. When Microsoft patched the vulnerability in May, it warned that computers that failed to install the fix could suffer a similar fate if reliable attack code ever becomes available. The reason: like the flaw that EternalBlue exploited, BlueKeep allowed for self-replicating attacks. Like a falling line of dominoes, a single exploit could spread from vulnerable machine to vulnerable machine with no interaction required of end users.

The risk was so great that Microsoft again implored customers to patch a month after its release. NSA officials also urged people to install the fix.

A big deal

As noted earlier, the module Metasploit developers released on Friday isn't quite as advanced as the leaked EternalBlue exploit, but it's still pretty effective. And that comes as both good and bad news for people who defend systems against malicious hacks.

"The release of this exploit is a big deal because it will put a reliable exploit in the hands of both security professionals and malicious actors," Ryan Hanson, principal research consultant at Atredis Partners and a developer who helped work on the release, told Ars. "I'm hoping the exploit will be primarily used by offensive teams to demonstrate the importance of security patches, but we will likely see criminal groups modifying it to deliver ransomware as well."

He continued:

> *It's not very often that you see Microsoft release a warning like they did with this bug. I'm sure the warning caused defensive teams to be more diligent about ensuring that all vulnerable systems were patched quickly, which was the purpose of the warning. However, Microsoft's warning was more of a "Capture the Flag" challenge for those of us on the offensive side. I rarely reverse security patches, but I became very curious and decided to reverse the patch as a learning exercise and also to figure out why Microsoft considered this bug to be so dangerous. Only a few days after the patch, people started sharing proof they had already reversed the patch and triggered a crash. Not long after, proof of successful code execution was shared by multiple people, including myself.*
>
> *Although several people had publicly proven code execution, nobody released their PoCs, which I assume is because we all realized exactly why Microsoft warned everyone about the dangers of this bug. Shortly after people started showing*

> *proof of code execution, the NSA also released an advisory regarding the risks associated with BlueKeep. With all the warnings and risks associated with this bug, it is pretty significant that an exploit will be released publicly for the first time. Especially after so many researchers have kept their PoCs private.*

A single machine is all it takes

Another of the primary developers behind the release is Sean Dillon, a senior security researcher at RiskSense. Friday's release is almost identical to the BlueKeep exploit video he published in June. It showed the module connecting to an unpatched Windows Server 2008 R2 computer, and using the exploit, had highly privileged System privileges. Dillon then used the open source Mimikatz application to obtain the cryptographic hashes of passwords belonging to other computers on the same network the hacked machine was connected to.

The ability to dump credentials used to connect to other computers underscores a key danger posed by the vulnerability. A single vulnerable machine could be used to infect all other machines in a network even if they're fully up to date. Dillon's video graphically portrayed this threat in June. With the open source code now available for anyone to examine, rewrite, or repurpose, the risk will be even harder for people to ignore.

"As an open-source project, one of Metasploit's guiding principles is that knowledge is most powerful when shared," s Brent Cook, the Metasploit engineering manager at security firm Rapid7 wrote in a post published on Friday. "Democratic access to attacker capabilities, including exploits, is critical for defenders—particularly those who rely on open-source tooling to understand and effectively mitigate risk."

https://arstechnica.com/?p=1563639

Zero-day privilege escalation disclosed for Android

reader comments

46 with 30 posters participating

Share this story

Researchers have disclosed a zero-day vulnerability in the Android operating system that gives a major boost to attackers who already have a toe-hold on an affected device.

The privilege-escalation flaw is located in the V4L2 driver, which Android and other Linux-based OSes use to capture real-time video. The vulnerability results from a "lack of validating the existence of an object prior to performing operations on the object," researchers with Trend Micro's Zero Day Initiative said in a blog post published Wednesday. Attackers who already have untrusted code running with low privileges on a device can exploit the bug to access privileged parts of the Android kernel. The severity score is rated a 7.8 out of a possible 10 points.

Modern OSes have become increasingly hard to compromise in recent years thanks to exploitation mitigations that prevent untrusted code from interacting with hard drives, kernels, and other sensitive resources. Hackers have responded by chaining two or more exploits together. A buffer overflow, for instance, may allow an attacker to load malicious code into

memory, and a privilege-escalation flaw gives the code the privileges it needs to install a persistent payload.

The net result of all of this: privilege-escalation bugs are increasingly valuable, as demonstrated by the so-called Dirty Cow vulnerability discovered affecting Linux in 2016. Within days of being discovered, the privilege-escalation bug was being used to root Android devices. A year after coming to light, Dirty Cow was being exploited by malicious apps to bypass security protections built into Android.

“This vulnerability is similar to Dirty Cow in that it is in the core code of the kernel, so it would apply to all Android devices,” Christoph Hebeisen, director of security intelligence at mobile security provider Lookout, told Ars. “However, an exploit based on this vulnerability would not be as elegant as DirtyCow and probably not quite as reliable.”

Based on the advisory, Hebeisen said it appears only apps or code that already have access to the V4L subsystem used by an attached camera could exploit the flaw. Dirty Cow, by contrast, resided in a core memory-management feature. What’s more, exploiting the flaw was relatively easy, a trait that made exploits highly reliable.

ZDI’s Wednesday post said researchers notified Google of the vulnerability in mid-March and that by the end of June, the company had confirmed that the flaw would be fixed. When ZDI asked Google for an update last month, Google responded there would be no further updates. Google released the Android Security Bulletin for September on Tuesday, and the flaw still wasn’t fixed. Google didn’t respond to a request for comment.

“Given the nature of the vulnerability, the only salient mitigation strategy is to restrict interaction with the service,” ZDI researchers wrote in Tuesday’s post. “Only the clients and servers that have a legitimate procedural relationship with the service should be permitted to

communicate with it.”

In an email, ZDI Director Brian Gorenc said Android users should limit the apps they install and ensure the apps come only from Google Play. He said ZDI verified the flaw affected the latest versions of Android.

https://arstechnica.com/?p=1562969

A glut of iOS 0-days pushes their price below cost of those for Android

reader comments

60 with 49 posters participating

Share this story

For the first time ever, the security exploit broker Zerodium is paying a higher price for zero-day attacks that target Android than it pays for comparable attacks targeting iOS.

An updated price list published Tuesday shows Zerodium will now pay $2.5 million apiece for “full chain (Zero-Click) with persistence” Android zero-days compared with $2 million for iOS zero-days that meet the same criteria. The previous program overview offered $2 million for unpublished iOS exploits but made no reference at all to the exploits for

Android. Zerodium founder and CEO Chaouki Bekrar told Ars the broker paid on a “case by case basis depending on the chain” for Android exploits.

“Flooded by iOS exploits”

Bekrar told Ars the move was prompted by a glut of working iOS exploit chains that has coincided with the growing difficulty of finding comparable exploits for versions 8 and 9 of Android. In a message, Bekrar wrote:

> *During the last few months, we have observed an increase in the number of iOS exploits, mostly Safari and iMessage chains, being developed and sold by researchers from all around the world. The zero-day market is so flooded by iOS exploits that we’ve recently started refusing some [of] them.*
>
> *On the other hand, Android security is improving with every new release of the OS thanks to the security teams of Google and Samsung, so it became very hard and time consuming to develop full chains of exploits for Android and it’s even harder to develop zero click exploits not requiring any user interaction.*
>
> *In accordance with these new technical challenges related to Android security and our observations of market trends, we believe that time has come to allocate the highest bounties to Android exploits until Apple re-improves the security of iOS and strengthens its weakest parts which are iMessage and Safari (Webkit and sandbox).*

Modern operating systems contain a variety of security protections that typically require attackers to combine two or more exploits in an attack chain, with each link tackling a different application or defense. Zero-click exploits are those that don’t require any interaction at all on the part of the end user. An exploit that arrives in a text message and

allows the attacker to take control of a device is an example. A one-click exploit, by contrast, requires the end user to take minimal action, such as visiting a booby-trapped website.

Wakeup call

The price change comes four days after researchers from Google's Project Zero reported that users of fully patched versions of iOS were vulnerable to iOS zero-days that were exploited in the wild for more than two years. Attacks against 14 separate vulnerabilities were packaged into five separate exploit chains that gave the attackers the ability to compromise up-to-date devices.

The attacks were waged from a small collection of hacked websites that used the exploits to indiscriminately attack every iOS device that visited. Attackers used the exploits to install malware that stole photos, emails, log-in credentials, live location data, and more from iPhones and iPads. Project Zero researchers didn't identify any of the websites that hosted the exploits. On Monday, researchers from security firm Volexity identified 11 websites serving Uyghur and East Turkistan visitors that likely served the iOS exploits. The Volexity post said one of the sites also appeared to exploit an Android vulnerability that stopped working in 2017 with the release of Chrome 60.

The Project Zero report that websites openly and indiscriminately exploited iOS zero-days for more than two years challenged many of the conventional assumptions some security researchers made about security on the Apple mobile OS. Previously, many assumed zero-click or one-click attack chains that worked against the latest version of iOS were so costly and rare that they were used sparingly. The haphazard way the exploits were used on the sites discovered by Project Zero suggested unpublished iOS attacks were plentiful, despite the considerable expertise needed to develop them.

“The latest set of zero-days affecting Apple’s platform announced by Google’s Project Zero were a bit of a wakeup call shattering our views on the iOS ecosystem and its security,” Jérôme Segura, director of threat intelligence at antivirus provider Malwarebytes, told Ars. “While it’s true that Apple controls the hardware and that OS updates are adopted quickly, we are seeing evidence that determined attackers are able to bypass iOS security mechanisms more than in the past.”

Zerodium’s update said the $2.5 million price applied to Android versions 8 and 9. The update made no reference to Android 10, which was released on Tuesday, but Bekrar told Ars that that version is covered as well. While Zerodium is paying $2.5 million and $2 million for zero-click exploit chains for Android and iOS, respectively, top price for comparable exploits targeting desktop OSes tops out at $1 million.

“Mobile users should not be worried as the overall security of mobile devices is nowadays much better than any laptop or computer,” Bekrar said.

https://arstechnica.com/?p=1561649

---

iOS vulnerability that let you jailbreak your iPhone is once again dead

reader comments

56 with 48 posters participating

Share this story

The iOS vulnerability that made it possible for users and hackers to jailbreak fully up-to-date iPhones and iPads is no more, following an update Apple released on Monday that patches the highly unusual bug.

The security advisory accompanying the release of iOS version 12.4.1 says it patches a kernel vulnerability that allows malicious apps to execute code that runs with the highest of privileges. (The use-after-free vulnerability was first fixed in iOS 12.3.) Then, last weekend, researchers noticed that version 12.4, released in June, was once again vulnerable. Jailbreak enthusiasts—who like the freedom that such vulnerabilities permit—quickly capitalized on the Apple developer mistake by releasing exploits that worked on fully patched devices.

Jailbreaking phones allows users to do all kinds of things that aren't normally possible, including installing unauthorized apps. But that freedom can potentially work to the favor of malicious hackers. Jailbreaking may weaken iOS protections that prevent an app from reading or modifying another app's data.

Jailbreaking can also allow malicious apps to silently install spyware. In 2016, a malicious attack known as Pegasus that was delivered by text message used a then-secret jailbreak vulnerability to install stealthy malware on iOS devices targeted by the United Arab Emirates government.

Over the years, publicly known jailbreaks have grown increasingly rare. The jailbreak discovered last weekend was all the more unusual because it was the result of a flaw Apple developers already knew of and, indeed, had already fixed.

That vulnerability—tracked as CVE-2019-8605—was originally found by Google researcher Ned Williamson. Somehow, the same flaw appears also to have been reintroduced into macOS version 10.14.6, according to an advisory that was published on Monday for a supplemental update of the OS.

Apple's Monday advisories credited a jailbreak researcher who goes by the handle Pwn20wnd with the discovery of the bug regression in both iOS and macOS. On Twitter, the researcher confirmed the mobile update effectively killed the jailbreak. While the researcher advised people to "Stay on iOS 12.4!" Ars recommends everyone update as soon as possible.

https://arstechnica.com/?p=1558201

Hackers are actively trying to steal passwords from two widely used VPNs

reader comments

9 with 9 posters participating

Share this story

Hackers are actively unleashing attacks that attempt to steal encryption keys, passwords, and other sensitive data from servers that have failed to apply critical fixes for two widely used virtual private network (VPN) products,

researchers said.

The vulnerabilities can be exploited by sending unpatched servers Web requests that contain a special sequence of characters, researchers at the Black Hat security conference in Las Vegas said earlier this month. The pre-authorization file-reading vulnerabilities resided in the Fortigate SSL VPN, installed on about 480,000 servers, and the competing Pulse Secure SSL VPN, installed on about 50,000 machines, researchers from Devcore Security Consulting reported.

The Devcore researchers discovered other critical vulnerabilities in both products. These make it possible for attackers to, among other things, remotely execute malicious code and change passwords. Patches for the Fortigate VPN became available in May and in April for Pulse Secure. But installing the patches can often cause service disruptions that prevent businesses from carrying out essential tasks.

Spraying the Internet

Over the past 36 hours, hackers have started spraying the Internet with code that attempts to opportunistically exploit that difficulty, independent researcher Kevin Beaumont said. He said he found attacks against Fortigate servers coming from 91.121.209.213, an IP address that has a history of previous misconduct. A scan on Friday using the BinaryEdge search engine showed a new IP address, 52.56.148.178, had also begun spraying exploits for the same vulnerability.

Earlier this month, two samples of exploit code for CVE-2018-13379, as the vulnerability is tracked, became publicly available here and here. The first one actually obtains data stored on vulnerable machines, while the latter merely checks if a machine is vulnerable.

Meanwhile, Beaumont said, attacks trying to exploit unpatched Pulse Secure servers are coming from 2.137.127.2. Exploit code became publicly available earlier this week. Independent researcher Troy Mursch said he also found attacks coming from 81.40.150.167 that also attempt to exploit or test for the vulnerability, which is indexed as CVE-2019-11510. In the event one of the mass scans identifies a vulnerable server, it may then exploit a code-execution flaw the Devcore researchers also discovered.

“These scans are targeting endpoints that are vulnerable to arbitrary file reading leading to sensitive information disclosure of private keys and user passwords,” Mursch told Ars. “They’re exploiting this vulnerability to read the contents of the `etc/passwd file to steal credentials. These credentials can then be used to conduct further command injection attacks (CVE-2019-11539) and gain access to the private network allowing for further malicious activity.”

Mursch said the honeypot server he used to detect the attacks was also able to identify that the IP address 2.137.127.2 was also targeting the Pulse Secure vulnerability. He said he didn’t believe either of the IPs was operated by a researcher who was merely scanning for unpatched servers. His honeypot was unable to detect code attacking the Fortigate vulnerability. Beaumont was using a honeypot provided by

BinaryEdge.

The vulnerabilities are serious because they affect a piece of software that’s required to be accessible to the Internet and that acts as a gateway to highly sensitive parts of an organization’s network. Obtaining hashed and in some cases plain-text passwords, encryption keys, and other sensitive data, could allow people to penetrate those networks. With more work, attackers who identify unpatched servers could also exploit the other vulnerabilities the Devcore researchers found. One Fortigate flaw, which they dubbed “The Magic Backdoor,” allows remote attackers who know a hard-coded key to change passwords.

Representatives from both Fortinet and Pulse Secure said the companies have been urging customers for months to patch their systems as soon as possible. Neither company could confirm or expand upon the reports of scanning coming from Beaumont and Mursch. Organizations using either of these VPNs should take time now to make sure they’re not vulnerable.

https://arstechnica.com/?p=1557239

Valve says turning away researcher reporting Steam vulnerability was a mistake

reader comments

60 with 39 posters participating

Share this story

In an attempt to quell a controversy that has raised the ire of white-hat hackers, the maker of the Steam online game platform said on Thursday it made a mistake when it turned away a researcher who recently reported two separate vulnerabilities.

In its statement, Valve Corporation references HackerOne, the reporting service that helps thousands of companies receive and respond to vulnerabilities in their software or hardware. The company also writes:

> *We are also aware that the researcher who discovered the bugs was incorrectly turned away through our HackerOne bug bounty program, where his report was classified as out of scope. This was a mistake.*
>
> *Our HackerOne program rules were intended only to exclude reports of Steam being instructed to launch previously installed malware on a user's machine as that local user. Instead, misinterpretation of the rules also led to the exclusion of a more serious attack that also performed local privilege escalation through Steam.*
>
> *We have updated our HackerOne program rules to explicitly state that these issues are in scope and should be reported. In the past two years, we have collaborated with and rewarded 263 security researchers in the community helping us identify and correct roughly 500 security issues, paying out over $675,000 in bounties. We look forward to continuing to work with the security community to improve the security of our products through the HackerOne program.*
>
> *In regards to the specific researchers, we are reviewing the*

> *details of each situation to determine the appropriate actions. We aren't going to discuss the details of each situation or the status of their accounts at this time.*

Valve's new HackerOne program rules specifically provide that "any case that allows malware or compromised software to perform a privilege escalation through Steam, without providing administrative credentials or confirming a UAC dialog, is in scope. Any unauthorized modification of the privileged Steam Client Service is also in scope."

Kicking the hornet's nest

Vasily Kravets

The statement and the policy change from Valve came two days after security researcher Vasily Kravets, an independent researcher from Moscow, received an email telling him that Valve's security team would no longer receive his vulnerability reports through the HackerOne bug-reporting service. Valve turned Kravets away after he reported a Steam vulnerability that allowed hackers who already had a toe-hold on a vulnerable computer to burrow into privileged parts of an operating system. Valve initially told Kravets such vulnerabilities were out of scope and gave no indication that the one Kravets reported would be fixed.

Valve's response rankled hackers and security professionals because so-called privilege-escalation vulnerabilities are something that Google, Microsoft, and mature open source developers routinely and readily fix in their products. Valve's contention that a demonstrated flaw of this type wasn't a legitimate vulnerability ran counter to long-standing

security norms. As criticism mounted, Valve quietly issued a patch, but researchers found that it could be bypassed. To make matters worse, Kravets on Tuesday publicly disclosed a new privilege escalation vulnerability in Steam. Valve's Thursday statement said both vulnerabilities reported by Kravets have now been fixed.

Another researcher is shut down

Criticism grew more fevered—and began to increasingly be directed at HackerOne—after Matt Nelson, a second independent researcher, said he used the vulnerability reporting service in June to disclose one of the same Steam flaws Kravets found. According to an exchange Nelson posted, a HackerOne representative said the vulnerability was out of scope to qualify for Valve's bug bounty program. When Nelson replied that he wasn't looking for money but instead only wanted the public to be aware of the vulnerability, the HackerOne representative asked Nelson to "please familiarize yourself with our disclosure guidelines and ensure that you're not putting the company or yourself at risk. https://www.hackerone.com/disclosure-guidelines."

Those guidelines specifically state, "Please note that we will not consent to disclose reports if they have been marked out-of-scope or inapplicable, or where Valve has not taken a specific corrective action / mitigation."

The HackerOne representative also locked the thread so that it no longer could be read. Nelson told Ars he was surprised by HackerOne's response. Using the abbreviation for non-disclosure agreement, he wrote in an email:

> *I wouldn't go as far as calling it an NDA, but they certainly were implying there could be repercussions for violating Valve's policy or HackerOne's disclosure policy. My reaction is that I could care less. If they want to ban a researcher submitting bugs in good faith, they certainly can. I'm not*

> *reporting via HackerOne to make money, it simply seemed like the easiest way to get in contact with them. Making the process more difficult for those willing to report vulnerabilities does absolutely nothing but hurt them and the general public.*

Nelson then reported the vulnerability directly to Valve. Valve, he said, acknowledged the report and "noted that I shouldn't expect any further communication." He never heard anything more from the company.

Shifting the blame

Nelson said he welcomed Valve's admission that it was a mistake for Valve to have turned away Kravets and the policy change that makes privilege-escalation vulnerabilities in scope.

"I can certainly believe that the scoping was misinterpreted by HackerOne staff during the triage efforts," Nelson wrote. "It is mind-blowing to me that the people at HackerOne who are responsible for triaging vulnerability reports for a company as large as Valve didn't see the importance of Local Privilege Escalation and simply wrote the entire report off due to misreading the scope."

In a post reporting the first vulnerability, Kravets also said that HackerOne staff told him he wasn't permitted to publicly disclose the vulnerability.

Meanwhile, he also told Ars on Thursday morning that he had yet to receive any communication from Valve and that he remained locked out of the Valve bug-reporting section of HackerOne.

Vasily Kravets
A HackerOne spokeswoman told Ars:

> *We aim to explicitly communicate our policies and values in all cases and here we could have done better. Vulnerability disclosure is an inherently murky process and we are, and have always been, committed to protecting the interests of hackers.*
>
> *Our disclosure guidelines emphasize mutual respect and empathy, encouraging all to act in good faith and for the benefit of the common good.*

Still cause for concern

While Valve's admission and policy change are encouraging, HackerOne's barring of Kravets remains an ongoing concern. HackerOne helps thousands of organizations receive and respond to vulnerability reports. Its policies have a huge effect on the security of products and devices used all over the world, said Katie Moussouris, founder and CEO of Luta Security who wrote the disclosure policies here and here, adopted by the International Organization for Standardization, while at Microsoft.

"Silencing the researcher on one issue is in complete violation of the ISO standard practices, and banning them from reporting further issues is simply irresponsible to affected users who would otherwise have benefited from these

researchers continuing to engage and report issues privately to get them fixed," Moussouris told Ars. "The norms of vulnerability disclosure are being warped by platforms that put profits before people."

This post was updated to remove claims Valve disputed the vulnerability in the official CVE advisory.

https://arstechnica.com/?p=1556407

Four wormable bugs in newer versions of Windows need your attention now

reader comments

39 with 30 posters participating

Share this story

Microsoft is warning of a four new Windows vulnerabilities that are "wormable," meaning they can be exploited to spread malware from one vulnerable computer to another without any user action in much the way the self-replicating WannaCry and NotPetya outbreaks did in 2017.

Similar to the so-called BlueKeep vulnerability Microsoft patched in May, the four bugs the company patched on Tuesday reside in Remote Desktop Services (RDS), which allow a user to

take control of a remote computer or virtual machine over a network connection. The bugs—indexed as CVE-2019-1181, CVE-2019-1182, CVE-2019-1222, and CVE-2019-1226—make it possible for unauthenticated attackers to execute malicious code by sending a specially crafted message when a protection known as Network Level Authentication is turned off, as is often done in large organizations.

In such networks, it's possible for exploits to ricochet from computer to computer. Leaving NLA on makes it harder for attacks to spread, since attackers must first have network credentials. The growing use of hacking tools such as Mimikatz, however, often enables attackers to surreptitiously obtain the needed credentials.

The race begins

Unlike BlueKeep—which affected only unsupported Windows versions or versions close to being unsupported—the bugs disclosed on Tuesday affect newer versions, specifically Windows 7, 8, and 10 and Server 2008, 2012, 2016, and 2019. That puts a much larger and potentially more sensitive fleet of computers at risk. Microsoft rated the severity of the vulnerabilities as 9.7 and 9.8 out of a possible 10. The company also said the chances of in-the-wild exploitation are "more likely."

"The vulnerabilities include the latest versions of Windows, not just older versions like in BlueKeep," independent security researcher Kevin Beaumont told Ars. "There will be a race between organizations to patch systems before people reverse engineer the vulnerability from the patches to learn how to exploit them. My message would be: keep calm and patch."

Windows machines that have automatic updating enabled should receive the patch within hours if they haven't already. Installing Tuesday's patches is the single most effective way

to ensure computers and the networks they're connected to are safe against worms that exploit the newly described vulnerabilities. For people or organizations that can't update immediately, a good mitigation is to "enable NLA and leave it enabled for all external and internal systems," Beaumont said in a blog post.

Enabling NLA doesn't provide an absolute defense against attacks. As noted earlier, attackers who manage to obtain network credentials can still exploit the vulnerabilities to execute code of their choice. Still, turning on NLA significantly increases the requirement, since the exploits can completely bypass the authentication mechanism built into RDS itself.

Harden the RDS

According to a blog post published Tuesday by Director of Incident Response at the Microsoft Security Response Center Simon Pope, Microsoft researchers discovered the vulnerabilities on their own during a security review designed to harden the RDS. The exercise also led to Microsoft finding several less-severe vulnerabilities in RDS or the Remote Desktop Protocol (RDP) that's used to make RDS work. Pope said there's no evidence any of the vulnerabilities were known to a third party.

The exercise came three months after the patching of BlueKeep, which was reported to Microsoft by the UK's National Cyber Security Center. It's possible—although Pope gave no indication—that the review came in response to that tip from the NCSC.

Some security researchers have speculated the original source of BlueKeep vulnerability report was the Government Communications Headquarters, the UK's counterpart to the National Security Agency, as part of a vulnerabilities equity process that calls for bugs to be disclosed once their value

to national security has diminished.

“So it’ll be ironic if the GCHQ VEP killed a RDP bug because it only affect [sic] old boxes but then MS audited all of RDP and killed one of their goto new hotness bugs. (Another good reason not to kill bugs),” Dave Aitel, a former NSA hacker who now heads security firm Immunity wrote on Twitter.

> *So it’ll be ironic if the GCHQ VEP killed a RDP bug because it only affect old boxes but then MS audited all of RDP and killed one of their goto new hotness bugs. (Another good reason not to kill bugs)*
>
> *— daveaitel (@daveaitel) August 13, 2019*

Aitel later acknowledged the theory “may be totally crazy! :)”

Whatever the case, the four wormable bugs disclosed Tuesday represent a threat not just to the Internet but to the health care, shipping, transportation, and other industries that rely on it. Administrators and engineers would do well to devote as much time as necessary to research the vulnerabilities to ensure they aren’t exploited the way WannaCry and NotPetya were two years ago.

https://arstechnica.com/?p=1551093

Silent Windows update patched side channel that leaked data

from Intel CPUs

Enlarge / An uncovered Intel Core i5-3210M (BGA) inside of a laptop.

reader comments

43 with 34 posters participating, including story author

Share this story

Microsoft last month pushed a silent update that mitigated a serious vulnerability in all CPUs Intel has introduced since 2012, researchers who discovered the flaw said Tuesday.

The vulnerability—discovered and privately reported to Intel 12 months ago—resided in every CPU Intel has introduced since at least its Ivy Bridge line of processors and possibly earlier, a researcher from security firm Bitdefender told Ars. By abusing a performance capability known as speculative execution, attackers could open a side channel that leaks encryption keys, passwords, private conversations, and other secrets that are normally off limits.

The attack demonstrated in a research paper published by Bitdefender is similar to those disclosed in January 2018 under the names Spectre and Meltdown. Patches Microsoft released around the same time largely blunted those attacks. Microsoft's advisory described the flaw occurring "when certain central processing units (CPU) speculatively access memory," without mentioning Intel or any other chip maker. Bitdefender researchers, meanwhile, said they tested two AMD chips and found no evidence either was affected.

Bitdefender's researchers found that a chip instruction known as SWAPGS made it possible to revive the side channel, even on

systems that had the earlier mitigations installed. SWAPGS gets called when a computing event switches from a less-trusted userland function to a more sensitive kernel one. Proof-of-concept exploits developed by Bitdefender invoked the instruction to siphon contents normally restricted to kernel memory into user memory.

“What we have found is a way to exploit the SWAPGS instruction which switches from userland to kernel mode in such a way that we could… carry out a side-channel attack,” Bogdan Botezatu, Bitdefender’s director of threat research and reporting, told Ars. “By doing that, we are going to leak kernel memory into the user space even if there are security measures that should prevent us from doing that.”

Speculative execution attempts to make computers run faster by predicting instructions before they’re actually given by an application or operating system. When a predicted instruction doesn’t actually come to be, it’s supposed to be discarded as if it had no effect. The Spectre and Meltdown research, however, showed that the effects can still be found in the lowest-level architectural features of the processor.

Speculatively executed instructions can load data into cache even though the instruction is later aborted. That data can then be detected and inferred by other processes, since accessing it will be slightly faster than if it wasn’t cached. Other data in the processor, including predictions about which instruction branch should be followed, can similarly be probed for clues about sensitive information.

The vulnerability found by Bitdefender affects all Intel CPUs that support SWAPGS and another instruction called WSGRBASE. Bitdefender researchers said, at a minimum, this includes all processor lines since Ivy Bridge. But Botezatu said it’s possible earlier chips are also affected. Bitdefender was able to exploit the side channel when chips ran Windows. Botezatu said that, while the vulnerability technically exists when

affected chips run on other operating systems, it was "unfeasible" to exploit chips running Linux, Unix, FreeBSD, or macOS.

Exploiting the vulnerability using JavaScript is not possible, so that makes website drive-by attacks unfeasible as well.

Microsoft silently patched the vulnerability during last month's update Tuesday. Microsoft said the fix works by changing how the CPU speculatively accesses memory. The fix doesn't require a microcode update from computer manufacturers. The vulnerability is tracked as CVE-2019-1125.

A Microsoft representative wrote in a statement:

> *We're aware of this industry-wide issue and have been working closely with affected chip manufacturers and industry partners to develop and test mitigations to protect our customers. We released security updates in July, and customers who have Windows Update enabled and applied the security updates are protected automatically.*

An Intel representative wrote: "Intel, along with industry partners, determined the issue was better addressed at the software level and connected the researchers to Microsoft. It takes the ecosystem working together to collectively keep products and data more secure, and this issue is being coordinated by Microsoft."

Red Hat said in a statement that the vulnerability affected both Intel and AMD chips and urged users to update systems "as soon as errata are available." Bitdefender researchers, for their part, said they tested two AMD processors and "neither exhibited speculative behavior for the SWAPGS instruction." In a statement, AMD officials wrote:

> *AMD is aware of new research claiming new speculative execution attacks that may allow access to privileged kernel*

> *data. Based on external and internal analysis, AMD believes it is not vulnerable to the SWAPGS variant attacks because AMD products are designed not to speculate on the new GS value following a speculative SWAPGS. For the attack that is not a SWAPGS variant, the mitigation is to implement our existing recommendations for Spectre variant 1.*

Bitdefender researchers described two attack scenarios. The first is when SWAPGS isn't being executed speculatively when it should, and the second is just the opposite, when SWAPGS is being executed speculatively and should not be.

Each scenario had two variants: (1) when the attacker detects a value located at a specific kernel address and (2) when the attacker finds a value at a random kernel address. AMD's statement indicates that its processors are affected only by the second variant of the second scenario. The statement also indicates that existing Spectre mitigations are effective in this case.

The most likely scenario for exploitation would be against a cloud service, most likely by hackers working for a nation-state. The vulnerability makes it possible for one virtual machine to steal secrets residing inside of another virtual machine running on the same vulnerable CPU.

The exploit would "make sense for a state-sponsored attacker that has access to resources in a multi-tenant environment," Botezatu said. Even in that scenario, it would take hours for secrets to leak out. Still, he said, running the exploit over periods as long as a year might be possible, since there's no easy way to detect it.

The Bitdefender paper said researchers first reported the vulnerability to Intel 12 months ago, on August 7, 2018. Intel responded three weeks later by saying it already knew of the vulnerability and had no plans to fix it. Bitdefender said it spent the next eight months insisting to Intel that the

behavior was problematic. Intel finally confirmed the leak of kernel memory on April 2 and indicated that a fix would come from fixes in operating systems.

Microsoft confirmed to Bitdefender it was investigating the leak on April 17. Bitdefender provided a new proof-of-concept exploit on April 22. On May 7, Microsoft reported it was able to reproduce the leak and would attempt to release a patch in July.

While the vulnerability isn’t likely to be widely exploited—if at all—it’s a testament to the difficulty of completely patching a new class of CPU flaws that stem from speculative execution. Since Spectre was disclosed 19 months ago, researchers have unearthed a raft of related ones. Don’t be surprised if more follow in the coming months or years.

https://arstechnica.com/?p=1547147