


Continue



Broadcom warned customers today about three VMware zero-days, tagged as exploited in attacks and reported by the Microsoft Threat Intelligence Center. The vulnerabilities (CVE-2025-22224, CVE-2025-22225, and CVE-2025-22226) impact VMware ESX products, including VMware ESXi, vSphere, Workstation, Fusion, Cloud Foundation, and Telco Cloud Platform. Attackers with privileged administrator or root access can chain these flaws to escape the virtual machine's sandbox. "This is a situation where an attacker who has already compromised a virtual machine's guest OS and gained privileged access (administrator or root) could move into the hypervisor itself," the company explained today. "Broadcom has information to suggest that exploitation of these issues has occurred 'in the wild'." Broadcom says CVE-2025-22224 is a critical-severity VCM heap overflow vulnerability that enables local attackers with administrative privileges on the targeted VM to execute code as the VMX process running on the host. CVE-2025-22225 is an ESXi arbitrary write vulnerability that allows the VMX process to trigger arbitrary kernel writes, leading to a sandbox escape, while CVE-2025-22226 is described as an HGS information-disclosure flaw that lets threat actors with admin permissions to leak memory from the VMX process. When asked if Microsoft had more information regarding the flaws 'in-the-wild' exploitation, a spokesperson told BleepingComputer the company "does not have anything additional to share at this time." VMware vulnerabilities are often targeted in attacks by ransomware gangs and state-sponsored hacking groups because they are commonly used in enterprise operations to store or transfer sensitive corporate data. Most recently, Broadcom warned in November that attackers were actively exploiting two VMware vCenter Server vulnerabilities that were patched in September. One allows privilege escalation to root (CVE-2024-38813) while the other is a critical remote code execution flaw (CVE-2024-38812) reported during China's 2024 Matrix Cup hacking contest. In January 2024, Broadcom also revealed that Chinese state hackers had exploited a critical vCenter Server vulnerability (CVE-2023-34048) as a zero-day since at least late 2021 to deploy VirtualPita and VirtualPie backdoors on vulnerable ESXi hosts. Patching used to mean complex scripts, long hours, and endless fire drills. Not anymore. In this new guide, Times breaks down how modern IT orgs are leveling up with automation. Patch faster, reduce overhead, and focus on strategic work -- no complex scripts required. Cisco has released patches to address three vulnerabilities with public exploit code in its Identity Services Engine (ISE) and Customer Collaboration Platform (CCP) solutions. The most severe of the three is a critical static credential vulnerability tracked as CVE-2025-20286, found by GMO Cybersecurity's Kentaro Kawane in Cisco ISE. This identity-based policy enforcement software provides endpoint access control and network device administration in enterprise environments. The vulnerability is due to improperly generated credentials when deploying Cisco ISE on cloud platforms, resulting in shared credentials across different deployments. Unauthenticated attackers can exploit it by extracting user credentials from Cisco ISE cloud deployments and using them to access installations in other cloud environments. However, as Cisco explained, threat actors can exploit this flaw successfully only if the Primary Administration node is deployed in the cloud. "A vulnerability in Amazon Web Services (AWS), Microsoft Azure, and Oracle Cloud Infrastructure (OCI) cloud deployments of Cisco Identity Services Engine (ISE) could allow an unauthenticated, remote attacker to access sensitive data, execute limited administrative operations, modify system configurations, or disrupt services within the impacted systems," the company explained. "The Cisco PSIRT is aware that proof-of-concept exploit code is available for the vulnerability that is described in this advisory." Cisco added that the following ISE deployments are not vulnerable to attacks: All on-premises deployments with any form factors where artifacts are installed from the Cisco Software Download Center (ISO or OVA). This includes appliances and virtual machines with different form factors. ISE on Azure VMware Solution (AVS) ISE on Google Cloud VMware Engine ISE on VMware cloud in AWS ISE hybrid deployments with all ISE Administrator personas (Primary and Secondary Administration) on-premises with other personas in the cloud. The company advises admins still waiting for a hotfix or who cannot immediately apply the hotfixes released today to run the application reset-config ise command on the Primary Administration persona cloud node to reset user passwords to a new value. However, admins should also be aware that this command will reset Cisco ISE to the factory configuration and that restoring backups will also restore the original credentials. The other two security bugs with proof-of-concept exploit code patched today are an arbitrary file upload (CVE-2025-20130) in Cisco ISE and an information disclosure (CVE-2025-20129) in the Cisco Customer Collaboration Platform (formerly Cisco SocialMiner). In September, Cisco patched another ISE flaw, a command injection vulnerability with public exploit code that can let attackers escalate privileges to root on unpatched systems. Patching used to mean complex scripts, long hours, and endless fire drills. Not anymore. In this new guide, Times breaks down how modern IT orgs are leveling up with automation. Patch faster, reduce overhead, and focus on strategic work -- no complex scripts required. Qualcomm has released security patches for three zero-day vulnerabilities in the Adreno Graphics Processing Unit (GPU) driver that impact dozens of chipsets and are actively exploited in targeted attacks. The company says two critical flaws (tracked as CVE-2025-21479 and CVE-2025-21480) were reported through the Google Android Security team in late January, and a third high-severity vulnerability (CVE-2025-27038) was reported in March. The first two are both Graphics framework incorrect authorization weaknesses that can lead to memory corruption because of unauthorized command execution in the GPU micronode while executing a specific sequence of commands, while CVE-2025-27038 is a use-after-free causing memory corruption while rendering graphics using Adreno GPU drivers in Chrome. "There are indications from Google Threat Analysis Group that CVE-2025-21479, CVE-2025-21480, CVE-2025-27038 may be under limited, targeted exploitation," Qualcomm warned in a Monday advisory. "Patches for the issues affecting the Adreno Graphics Processing Unit (GPU) driver have been made available to OEMs in May together with a strong recommendation to deploy the update on affected devices as soon as possible." This month, Qualcomm has also addressed a buffer over-read in Data Network Stack & Connectivity (CVE-2024-53026) that unauthenticated attackers can exploit to gain access to restricted information using invalid RTPC packets sent during a VoLTE/WiFi IMS calls. In October, the company fixed another zero-day (CVE-2024-43047) that the Serbian Security Information Agency (BIA) and the Serbian police exploited to unlock seized Android devices belonging to activists, journalists, and protestors using Cellebrite's data extraction software. While investigating the attacks, Google's Threat Analysis Group (TAG) found evidence suggesting that devices were also infected with NoviSpy spyware using an exploit chain to circumvent Android's security mechanisms and install itself persistently at the kernel level. One year earlier, Qualcomm also warned that threat actors were exploiting three more zero-day vulnerabilities in its GPU and Compute DSP drivers. In recent years, the company has patched various other chipset security flaws that could let attackers access users' text messages, call history, media files, and real-time conversations. Patching used to mean complex scripts, long hours, and endless fire drills. Not anymore. In this new guide, Times breaks down how modern IT orgs are leveling up with automation. Patch faster, reduce overhead, and focus on strategic work -- no complex scripts required. With the recent attacks on UK retailers Marks & Spencer and Co-op, so-called Scattered Spider has been all over the media, with coverage spilling over into the mainstream news due to the severity of the disruption -- currently looking like hundreds of millions in lost profits for M&S alone. This coverage is extremely valuable for the cyber security community as it raises awareness of the battles that security teams are fighting every day. But it's also created a lot of noise that can make it tricky to understand the big picture. So here's three things that you might have missed -- some you probably know already, and others that you might not be aware of if you haven't been tracking Scattered Spider beyond the recent attacks. 1. There's no such thing as Scattered Spider as a community, we sometimes forget that giving cool names to patterns of threat actor activity can sensationalize and make supervillains out of criminals. That said, cool names are sticky, and have a better chance of being commonly recognized and adopted, which is helpful for intelligence sharing. But we need to remember that Scattered Spider didn't call themselves Scattered Spider, CrowdStrike did. And there are lots of other names given to the patterns of activity and techniques that we know as Scattered Spider: UNC3944 (Mandiant) Octo Tempest (Microsoft) Oktapus (Group-IB) Muddled Libra (Unit 42) Scatter Swine (Okta) But it's not quite as simple as that, because there aren't clear boundaries. The pattern of activity that analysts classify as Scattered Spider touches on a number of self-named criminal groups like, Lapsus\$, Yanluowang, Karakurt, and ShinyHunters (behind the Snowflake attacks in 2024). Typically, the main "brands" created by attackers overlap with ransomware/extortion crews, which often have their own unique (or at least modified) ransomware encryptor and platform. This explains the other cool name that's cropped up a lot in the recent reporting -- DragonForce -- creating some confusion around specifically who executed the attacks on M&S and Co-op. Unlike Scattered Spider, DragonForce is a Ransomware-as-a-Service group that provides tooling and specialist services for hire to affiliates like Scattered Spider. They are not the ones executing the attack, but the criminals classified under "Scattered Spider" are effectively using their services and encryption software once they have completed the initial intrusion. What defines Scattered Spider? So, it's confusing, but what we're really tracking is patterns of behavior tied to certain regions of operation. When you think of Scattered Spider, you might be reminded of the series of arrests that happened throughout 2024. And yet, attacks have continued -- because we're not talking about a tight-knit group of specific individuals, but a broader community or collective of criminals, all using similar techniques, with the same ultimate goal -- making money (typically through data theft, ransomware, and extortion). So, what defines so-called Scattered Spider? Primarily English native speakers located mainly in English-speaking countries -- the UK, US, Canada, Australia -- but with activity also traced to mainland Europe, Russia, and India. Scattered Spider presenceSource: Mandiant Use of predominantly identity-based tactics, techniques and procedures (TTPs) specialising in phishing, credential attacks, help desk scams/fishing, SIM swapping, smishing, etc. -- all designed to achieve account takeover. Cloud-conscious techniques, such as targeting modern cloud identity provider accounts such as Okta and Microsoft Entra, and abusing cloud services and environments. When we think of Scattered Spider, we think of the quintessential cloud-native attacker who has grown up in the modern era of computing and internet services where being a hacker is less about network exploits than it is about logging into accounts on apps and services. These are people who probably cut their teeth in credit card scams and other forms of internet fraud rather than trawling the internet for exposed servers and open ports. So they're identity-first, but more important than that, they're flexible and adaptable. They're also willing to go after any and every company that presents an opportunity. Scattered Spider has dominated the news in recent weeks following high-profile breaches impacting UK retailers. But with a long history of high-profile ransomware attacks, this is just one example of their identity-based approach. Join Push Security as they go beyond the breaches and learn how to defend your organization against Scattered Spider's growing arsenal of TTPs. Watch the webinar on-demand 2. Help desk scams aren't new The headline story from the recent campaign against UK retailers is the use of help desk scams. This typically involves the attacker calling up a company's help desk with some level of information -- at minimum, PII that allows them to impersonate their victim, and sometimes a password, leaning heavily on their native English-speaking abilities to trick the help desk operator into giving them access to a user account. How it works The goal of a help desk scam is to get the help desk operator to reset the credentials and/or MFA used to access an account so the attacker can take control of it. They'll use a variety of backstories and tactics to get that done, but most of the time it's as simple as saying "I've got a new phone, can you remove my existing MFA and allow me to enroll a new one?" From there, the attacker is then sent an MFA reset link via email or SMS. Usually, this would be sent to, for example, a number on file -- but at this point, the attacker has already established trust and bypassed the help desk process to a degree. So asking "can you send it to this email address" or "I've actually got a new number too, can you send it to..." gets this sent directly to the attacker. At this point, it's simply a case of using the self-service password reset functionality for Okta or Entra (which you can get around because you now have the MFA factor to verify yourself) and voila, the attacker has taken control of the account. And the best part? Most help desks have the same process for every account -- it doesn't matter who you're impersonating or which account you're trying to reset. So, attackers are specifically targeting accounts likely to have top tier admin privileges -- meaning once they get in, progressing the attack is trivial and much of the typical privilege escalation and lateral movement is removed from the attack path. So, help desk scams have proved to be a reliable way of bypassing MFA and achieving account takeover -- the foothold from which to launch the rest of an attack, such as stealing data, deploying ransomware, etc. This isn't their first rodeo But something that's not quite coming across in the reporting is that Scattered Spider have been doing this successfully since 2022, with the M&S and Co-op attacks merely the tip of the iceberg. Fishing (calling a user to get them to give up their MFA code) has been a part of their toolkit since the beginning, with the early attacks on Twilio, LastPass, Riot Games, and Coinbase involving some form of voice-based social engineering. Notably, the high-profile attacks on Caesars, MGM Resorts, and Transport for London all involved calling a help desk to reset credentials as the initial access vector. Caesars in August 2023 where hackers impersonated an IT user and convinced an outsourced help desk to reset credentials, after which the attacker stole the customer loyalty program database and secured a \$15m ransom payment. MGM Resorts in September 2023, where the hacker used LinkedIn information to impersonate an employee and reset the employee's credentials, resulting in a 6TB data theft. After MGM refused to pay, the attack eventually resulted in a 36-hour outage, a \$100m hit, and a class-action lawsuit settled for \$45m. Transport for London in September 2024 resulted in 5,000 users' bank details exposed, 30,000 staff required to attend in-person appointments to verify their identities and reset passwords, and significant disruption to online services lasting for months. So not only have Scattered Spider been using these techniques for some time, but the severity and impact of these attacks has been ramping up. Avoiding help desk gotchas There's lots of advice for securing help desks being circulated, but much of the advice still results in a process that is either phishable or difficult to implement. Ultimately, organizations need to be prepared to introduce friction to their help desk process and either delay or deny requests in situations where there's significant risk. So, for example, having a process for MFA reset that recognizes the risk associated with resetting a high-privileged account. Require multi-party approval / escalation for admin-level account resets Require in-person verification if the process can't be followed remotely Freeze self-service resets when suspicious behavior is encountered (this would require some kind of internal process and awareness training to raise the alarm if an attack is suspected) And watch out for these gotchas: If you receive a call, good practice is to terminate the call and dial the number on file for the employee. But, in a world of SIM swapping, this isn't a foolproof solution -- you could just be re-dialing the attacker. If your solution is to get the employee on camera, increasingly sophisticated deepfakes can thwart this approach. But, help desks are a target for a reason. They're "helpful" by nature. This is usually reflected in how they're operated and performance measured -- delays won't help you to hit those SLAs! Ultimately, a process only works if employees are willing to adhere to it -- and can't be socially engineered to break it. Help desks that are removed from day-to-day operations (especially when outsourced or offshored) are also inherently susceptible to attacks where employees are impersonated. But, the attacks we're experiencing at the moment should give security stakeholders plenty of ammunition as to why help desk reforms are vital to securing the business (and what can happen if you don't make changes). 3. Scattered Spider don't just do help desk scams All that said, there's a bigger picture here -- help desk scams aren't the only tool in the Scattered Spider toolkit. They've consistently used a range of techniques, with a particular affinity for SIM swapping, smishing, and even basic credential phishing (usually targeted at Okta accounts). And this year, security researchers have observed Scattered Spider increasingly using Attacker-in-the-Middle (AiTM) phishing toolkits to bypass MFA. Scattered Spider phishing pages running EvilginxSource: Researchers at SilentPush This is very much on-brand for Scattered Spider. They exclusively use identity-based methods for their initial intrusions, all of which are designed to bypass MFA and achieve account takeover. Their attacks are usually very direct. Scattered Spider tend to go straight for accounts that have elevated permissions, enabling them to quickly progress their attack. For example, in the 2023 MGM attack, the attacker directly accessed an account with Super Admin permissions in Okta, which they combined with an inbound federation attack to impersonate any user in the tenant, get Azure admin privileges, and authenticate to the Azure-hosted VMware environment where they deployed ransomware. They've also demonstrated that they are specifically targeting VMware servers as their target for ransomware deployment/encryption, noted in the MGM and M&S attacks. By targeting the VMware hypervisor (usually by adding their compromised identity to the Admins group in VCentre), they're able to consciously evade endpoint-level controls running on the virtual machines themselves, such as EDR. Particularly if we consider the bigger picture with adjacent groups like ShinyHunters, who were behind the Snowflake attacks in 2024, and the severity of their attacks, we can see similar goals but different ways of achieving those goals. The Snowflake attacks leveraged stolen credentials from prior infostealer infections dating back to 2021 to log into accounts without MFA (with widespread MFA gaps a big problem due to the nature of Snowflake identity management at the time), resulting in hundreds of millions of breached records across 165 victims. You can also look at groups like Lapsus\$, who've demonstrated strikingly similar techniques in the past too. So in summary, Scattered Spider uses a range of identity-based techniques to take over privileged accounts for their initial intrusion, all of which are designed to bypass MFA. They aren't wedded to any specific technique though, and will use whatever means necessary within that identity-based framework to get the job done. Conclusion You can think of Scattered Spider as a kind of "post-MFA" threat actor that does everything they can to evade established security controls. By targeting identities and account takeovers, they bypass endpoint and network surfaces as much as possible, until the very end of the attack chain -- by which point it's almost too late to be relying on those controls. So, don't over-index on help desk scams -- you need to consider your broader identity attack surface and various intrusion methods, from apps and accounts with MFA gaps, local accounts giving attackers a backdoor into accounts otherwise accessed with SSO, and MFA-bypassing AiTM phishing kits that are the new normal for phishing attacks. Watch this on-demand webinar from researchers at Push Security to learn more about Scattered Spider's TTP evolution and what you can do to defend your organization. Learn how Push Security stops identity attacks Push Security provides comprehensive identity attack detection and response capabilities against techniques like AiTM phishing, credential stuffing, password spraying and session hijacking using stolen session tokens. You can also use Push to find and fix identity vulnerabilities across every app that your employees use, like: ghost logins; SSO coverage gaps; MFA gaps; weak, breached and reused passwords; risky OAuth integrations; and more. If you want to learn more about how Push helps you to detect and defeat common identity attack techniques, book some time with one of our team for a live demo. Sponsored and written by Push Security.