

eSENTIRE

vmware® Carbon Black

Detecting and Responding to Zero-Day Attacks

Protecting yourself when patching is not an option.




Table of Contents

3	Executive Summary
5	Introduction to Zero-Day Attacks
5	Discovering a vulnerability
5	From vulnerability to exploitation
6	Impact
6	Prevalence
8	Detecting Attacks on Day Zero (and Beyond)
8	ProxyLogon
8	Timeline
10	Detecting post-exploitation intrusion actions
12	Conclusions

Executive Summary

A zero-day attack is an attack that exploits a software vulnerability that is not known to the software vendor or its users. Because the vulnerability is not known to the vendor, there is no patch to repair it. Since the vulnerability is not known to the software's users, no specific protective actions are taken.

Consequently, a threat actor equipped with a zero-day exploit is presented with a land of opportunity. Attacks may proceed for weeks or months before investigations uncover the new exploit and vulnerability, and only then can the vendor begin to develop a patch. Until a patch is available, users are faced with an uneasy conundrum: keep using the vulnerable software or disable it until it has been repaired. In many cases, circumstances dictate that the software must remain in use.

Unfortunately, zero-day exploits are not particularly rare. Google's Project Zero tracks zero-day exploits observed in the wild (using a strict definition). Over the six years for which there is complete 12-month data, the group observed, on average, more than 22 zero-day exploits per year. Through the first four months of 2021, the group lists 18 in-the-wild exploits—on pace to significantly surpass the totals of previous years. And these are not vulnerabilities in niche products. 2021's list is dominated by Google, Microsoft, Apple and Adobe products relied upon by all manner of organizations, including the ProxyLogon vulnerabilities within Microsoft Exchange, which led to tens of thousands of organizations being at significant risk of compromise.

By definition, a novel zero-day attack cannot be prevented (e.g., by patching). In this context, an organization's ability to withstand a zero-day attack is completely dependent upon its capacity to detect and respond to an incident post-exploitation—this fact is true for a genuine zero-day exploit and for the period between an exploit becoming known and a patch becoming available. Of course, detection and response capabilities both require visibility across as much of the IT environment as is possible, including endpoints, networks and cloud environments.

Advanced protection platforms (e.g., endpoint protection, network protection) can employ behavioral analysis and machine learning to detect, to alert on and/or to automatically block, suspicious activity—thereby preventing an attack chain from proceeding, post-exploitation.

For threats that span technology boundaries and that are especially well-hidden, the ability to correlate and analyze vast volumes of data becomes the difference between detecting the threat in time to prevent damage or disruption and finding out only after the attacker has achieved their objectives. Extended detection and response (XDR) platforms that leverage machine learning can ingest endpoint, network, log, and cloud data sources to filter out noise and automatically block certain known and unknown threats.

These technological capabilities must be supported by a human-led threat response function to analyze the alerts, manually investigate and go on the hunt for other signs of an attacker's presence.

Once a zero-day becomes known publicly, an additional set of response activities should also be employed: vulnerability scanning and retroactive hunting for indicators of compromise (IoCs):

- Vulnerability scanning enables an organization to assess their degree of exposure to the new threat; in fact, vulnerability scanning is a general best-practice because it's not uncommon for patching programs to overlook resources or for unpatched systems to accidentally become externally exposed as a byproduct of configuration changes within the network
- Retro-hunting for IoCs (when such information is available) enables an organization to discover if it has already fallen victim

Zero-day attacks will always exist; therefore, it is imperative that every cybersecurity program consider what happens after the exploit has been used (e.g., to gain initial access, to escalate privileges, etc.), particularly the organization's capacity to quickly detect and respond to this post-exploitation behavior.

Introduction to Zero-Day Attacks

A zero-day attack is an attack that exploits a software vulnerability that is not known to the software vendor or its users. Because the vulnerability is not known to the vendor, there is no patch to repair it. Because the vulnerability is not known to the software's users, no specific protective actions are taken.

Consequently, a threat actor equipped with a zero-day exploit is presented with a land of opportunity

Discovering a vulnerability

Whether or not a newly discovered vulnerability leads to widespread exploitation is hugely dependent on who discovered it. If a well-intentioned researcher is the discoverer, the first public knowledge of the vulnerability may be the news that it is patched in a software update. If a threat actor is the discoverer, then the vendor and users may only become aware once it has already been exploited to devastating effect. The vulnerabilities targeted by zero-day exploits are discovered by:

- **Researchers within software vendors or organizations**, who strive to discover vulnerabilities so that they can be patched before others discover them
- **White hat researchers**, who typically notify vendors of vulnerabilities (so that they can be patched) before publishing the research at a later date
- **Nation state researchers**, whose goal is to discover vulnerabilities that pose a risk to national security (in which case they may notify vendors) or that can be exploited in the national interest
- **Black hat researchers**, whose primary aim is to develop exploits that can be used for malicious purposes

From vulnerability to exploitation

To take advantage of a vulnerability, a threat actor must have a way of exploiting it. Some vulnerabilities are easily exploitable, requiring only slight modifications to existing attack tools, while others are much more complex. With this latter group, there is often a series of incremental steps:

1. Vulnerability is discovered
2. Proof-of-concept (PoC) demonstrates the viability of exploitation
3. Exploit is weaponized in a practical technique
4. Exploit weaponization and tooling is optimized

Depending upon the nature of the vulnerability and the degree of difficulty in developing and weaponizing a viable exploit, this process typically plays out over anywhere from a few days to many months.

Once weaponized, an exploit may be used exclusively by a particular threat actor (at least until it's copied by another) or sold on dark web marketplaces.

CVE-2020-1472¹ (ZeroLogon) is an example that illustrates the path from vulnerability discovery to widespread exploitation. ZeroLogon was discovered by researchers at Secura and was announced on August 11, 2020, as part of the Microsoft August 2020 security updates patch. On September 11—after allowing a month for IT administrators to apply patches—Secura released technical details and a test tool.² The technical details publication set the clock ticking, as it served as a how-to guide for security researchers and would-be attackers.³

Just a few days later, PoC exploit code was available, the tooling rapidly evolved to become more effective, and attacks in the wild became imminent.

Impact

The impact of any particular zero-day exploit depends upon a number of factors, including:

- What the exploit achieves
- Where it fits in an attack chain
- How many vulnerable systems exist
- Who uses them

Continuing the previous example, ZeroLogon allows for privilege escalation,⁴ but the vulnerability is only exploitable after an attacker has already gained initial access⁵ into an environment.

Despite this restriction, ZeroLogon received a rare 10/10 under the Common Vulnerability Scoring System (CVSS),⁶ which incorporated the potential damage that can be wrought by an attacker with administrative privileges.

Notably, ZeroLogon is a vulnerability in Windows Server, which has an enormous and diverse install base. Owing to the significant risk associated with ZeroLogon, on Sept. 18 the Cybersecurity and Infrastructure Security Agency released a rare emergency directive requiring all U.S. federal agencies to update all Windows Servers with the domain controller role by 11:59 p.m. EDT, Monday, Sept. 21, 2020.⁷

Prevalence

Unfortunately, zero-day exploits are not particularly rare. Over the six complete years for which Google's Project Zero has provided tracking, the average is more than 22 zero-day exploits per year (Figure 1).⁸




Figure 1—Simplified attack workflow

It should be noted, too, that Google uses a strict definition of a zero-day vulnerability, omitting those that were opportunistically exploited by attackers in the gap between disclosure and a patch becoming available. Similarly, Zerologon does not feature in Google's tracking because it was disclosed to Microsoft long before it became public and was used in the wild.

Through the first four months of 2021, the group lists 18 in-the-wild exploits—on pace to significantly surpass the totals of previous years. And these are not vulnerabilities in niche products; 2021's list is dominated by Google, Microsoft, Apple and Adobe products relied upon by all manner of organizations.

This reality has important implications for cybersecurity as, by definition, zero-day attacks cannot be prevented (e.g., by patching). In this context, an organization's ability to withstand a zero-day attack is completely dependent upon its capacity to detect and respond to an incident post-exploitation.

Detecting Attacks on Day Zero (and Beyond)

Because a novel zero-day attack cannot be prevented—and even after a vulnerability becomes known, vulnerable systems may still be exposed until a patch is available—it is critical that an organization is able to detect and respond to whatever activities follow the initial exploitation.

The ProxyLogon saga illustrates the danger of zero-day exploits, but at the same time highlights the importance of advanced detection and response capabilities.

ProxyLogon

As of the end of April, 2021's most impactful zero-day exploitation is ProxyLogon, the name created by the researcher who discovered a collection of vulnerabilities in Microsoft Exchange Server versions 2010, 2013, 2016 and 2019.

Like the Citrix (CVE-2019-19781) episode of early 2020, ProxyLogon led to attackers achieving actions on objective in many instances.

ProxyLogon CVEs

ProxyLogon consists of the following Common Vulnerabilities and Exposures (CVEs):

- CVE-2021-26855 (CVSS Score: 9.1/10): A server-side request forgery (SSRF) vulnerability in Microsoft Exchange, which allows a threat actor to send an arbitrary HTTP request and authenticate as the Exchange server
- CVE-2021-26857 (CVSS Score: 7.8/10): An insecure deserialization vulnerability in the Unified Messaging service that requires administrator privileges or the use of another vulnerability to exploit; exploiting this vulnerability gives an attacker the controls to run code as SYSTEM on an Exchange server
- CVE-2021-26858/CVE-2021-27065 (CVSS Score: 7.8/10): Post-authentication arbitrary file write vulnerabilities in Microsoft Exchange; these require either compromising an administrator's credentials or the use of CVE-2021-26855

Timeline

When exploited, these vulnerabilities permit access to on-premises Exchange servers, thereby enabling unauthorized access to email. Attackers also employed web shell malware to maintain access to compromised Exchange servers.

ProxyLogon became known on March 2, 2021, when Microsoft released security updates detailing the vulnerabilities and also published a threat research report describing technical details related to these attacks.

At the time, Microsoft believed these attacks to be limited and targeted in nature, and attributed the activity to a threat actor group dubbed HAFNIUM. However, within hours it became clear that exploitation was much more widespread than Microsoft believed. This conclusion was possible because the publication of the technical details and the high-profile nature of the threat allowed cybersecurity researchers to investigate and caused organizations to dive into their Exchange deployments, both of which revealed far more attacks than were previously known.

Concerningly, the investigations found that attacks targeting these vulnerabilities had been ongoing since at least Jan. 2021 and indicated that all vulnerable servers (i.e., not just those of high-value targets) were at a high risk of exploitation by advanced threat actor groups.

In response to these further developments, IT personnel around the world rushed to review their Exchange servers for any evidence of exploitation.

By March 10, publicly available proof-of-concept exploit code and in-depth technical details for two of the vulnerabilities had been released. Prior to this development, only advanced threat actors—perhaps only HAFNIUM—were able to exploit these vulnerabilities, but the PoC code made exploitation possible for the wider cybercrime community, including less-skilled threat actors.

Widespread exploitation was now imminent—Shodan data from March 11 (Figure 2) showed tens of thousands of vulnerable hosts—and it seemed likely that these exploits would now be incorporated into the plethora of attack tools and the diverse range of services (e.g., by initial access brokers) available on the dark web.

ProxyLogon-Vulnerable Hosts by Country




Figure 2—Shodan data from March 11, 2021, showed more than 20,000 vulnerable hosts

True to expectation, ransomware attacks quickly followed the PoC release (Figure 3). For example:

- DearCry/DoejoCrypt incidents were observed in which attackers performed reconnaissance, installed Cobalt Strike and encrypted files
- BlackKingdom/Pydomer incidents featured credential dumping, reconnaissance, distribution of a Python payload and attempted file encryption and data theft




Figure 3—Attacks that leveraged ProxyLogon played out in three fairly distinct phases

Detecting post-exploitation intrusion actions

Figure 4 shows the ProxyLogon exploit chain that can lead to post-compromise hands-on-keyboard activity; while the zero-day exploitations themselves cannot be prevented until technical details are known, it remains possible in theory to detect post-exploitation activities (the green overlay) on day zero.

In practice, doing so requires state-of-the-art protection platforms. Such platforms are often able to detect much of the malicious activity that follows exploitation. However, for particularly novel techniques and tactics, the platforms may need to be updated as new information about the zero-day attack emerges.




Figure 4—The Exchange Server exploit chain [Source: Microsoft¹⁰]

Credential dumping, account discovery, reconnaissance, signed binary execution and coinminers are common signs of intrusion activity that are likely to take place after gaining initial access. Much of these activities can be performed without ingress of malicious tools, by using trusted Windows processes (living-off-the-land binaries, or LOLbins) to execute code and functions for malicious purposes. Because these actions are “fileless,” they often evade traditional anti-virus solutions.

One of the most common post-compromise activities observed following ProxyLogon exploitation was the deployment of web shells. Figure 5 shows that an attempt by the w3wp.exe process to launch an interactive web shell on an Internet Information Services (IIS) server was automatically detected and blocked by the Carbon Black Defense platform.

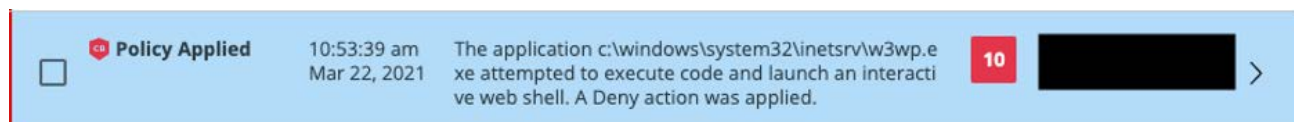


Figure 5—This attempt to execute code and launch an interactive web shell was automatically blocked.

While a significant amount of intrusion activity leverages processes that by themselves are legitimate, the manner in which they are used and spawned can be a sign of malicious activity. For example, Figure 6 shows that an attempt by the Outlook Web Access IIS Server Process to launch the w3wp.exe process was detected as a likely persistence action.



Figure 6—Carbon Black detects a suspicious child process and raises an alert

Attackers can leverage trusted processes to achieve a number of intrusion objectives. Figure 7 shows a process tree in which:

- verclsid.exe performed signed binary execution
- ntdsutil.exe attempted to dump OS credentials
- vssadmin.exe attempted to inhibit system recovery capabilities

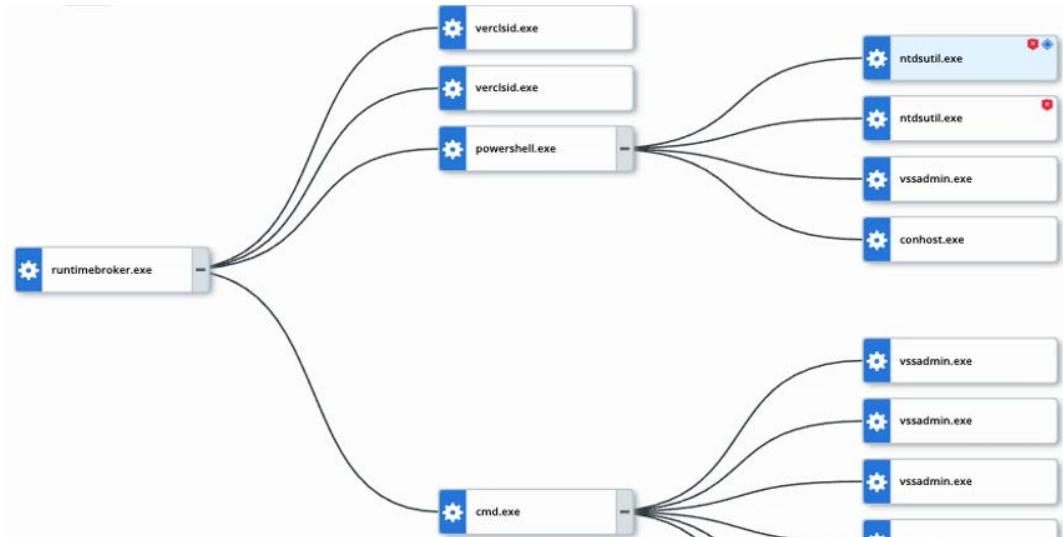


Figure 7—Many legitimate processes can be used with malicious intent.

In addition to the activities shown above, ProxyLogon exploitation was observed being followed by:

- Credential theft via Mimikatz, Vidar and RedLine
- Account discovery via net.exe
- Domain trust discovery via nltest.exe, adfind.exe and BloodHound
- Suspicious PowerShell, wscript (Windows Script host) and cscript (scripting languages can be used to execute any number of tools and Windows processes in order of preference, essentially automating large parts of intrusion reconnaissance)

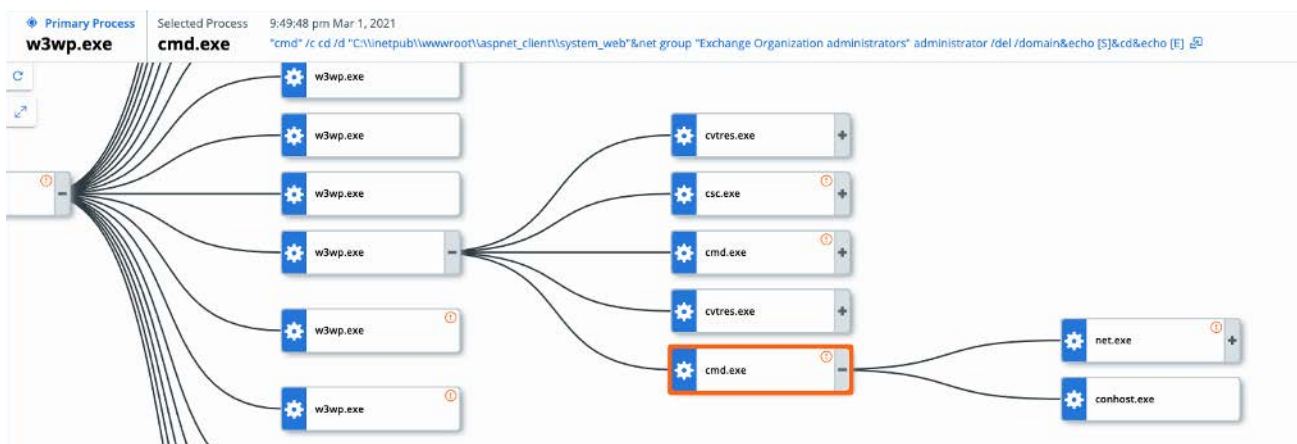


Figure 8—An attacker is detected trying to remove account access and delete the default Exchange administrator

Conclusions

Patching is a necessary element of any effective security strategy, but patching alone is insufficient.

Identifying a compromise requires visibility into your assets coupled with detection and response capabilities. While detecting a genuine zero-day exploit may be nearly impossible, threat actors rarely reinvent every aspect of an attack—that means it remains possible to detect known tradecraft post-exploit.

Vulnerability scanning

It goes without saying that as soon as a vendor releases a patch, it should be reviewed and—where possible—applied to remove the vulnerability.

However, history teaches that years-old vulnerabilities are still impacting organizations around the world. Recall that the theft of FireEye's Red Team tools was met with alarm,¹¹ despite the fact that many of the exploits leveraged were quite old.

The fact is that even a diligent patching program can be undermined by a number of factors, including:

- “Shadow IT” systems that operate without official knowledge
- Inadvertently re-exposing an unpatched system
- Inconsistent application of programs within distributed offices
- Legacy/forgotten systems that remain exposed

For example, one prominent theory suggests that the infamous Jones Day breach¹² was the result of a satellite office continuing to use a vulnerable (in fact, end-of-life) system, unbeknownst to the main firm—a risk that could be minimized with proactive vulnerability scanning.

Operationalizing information to bolster response

It's also important to understand that ensuring a strong defense against zero-day events demands more than just having advanced protection solutions in place ahead of time—it also requires operationalizing information as quickly as possible to continually bolster detection and response capabilities.

Within minutes of Microsoft making information available, eSentire's Managed Vulnerability Service (MVS) was updated to identify the ProxyLogon vulnerabilities.

During the height of ProxyLogon activity, findings from the research community came fast and furious. This information was rapidly operationalized¹³ and by March 4, eSentire had confirmed that:

- esENDPOINT identified suspicious Exchange processes and post-exploitation activity associated with known attacks
- eSentire's BlueSteel machine learning engine identifies malicious PowerShell activity associated with the esNETWORK identified exploitation of CVE-2021-26857 and CVE-2021-26855

By March 8, esLOG had been updated to identify exploitation of CVE-2021-26857 and CVE-2021-27065.

In parallel to the detection and automatic blocking activities touched on earlier, eSentire also actively reviewed customer networks for Indicators of Compromise (IoC), which would reveal if an attacker had successfully bypassed the defenses that were in place. While an organization can perform these hunts on their systems, doing so requires understanding the technical details of an attack and relating it to the way technology is integrated into the environment, as well as knowing how to effectively use logs, endpoint telemetry, anomaly detection, and so on—capabilities that may extend beyond the resourcing or skill set of IT teams that lack security specialists.

In addition to these retroactive hunts, eSentire's Threat Response Unit (TRU) examined Shodan data to identify any vulnerable customers and their IT providers.

Are you ready to respond to zero-day attacks?

Zero-day attacks will always be a threat, so it's imperative that your cybersecurity program consider what happens after an exploit has been used to gain entry into, and pursue intrusion actions within your environment.

If you're not currently engaged with a Managed Detection and Response provider, we highly recommend you partner with us for security services in order to disrupt threats before they impact your business.

If you're experiencing a security incident or breach contact us  1-866-579-2200

eSENTIRE

eSentire, Inc., is The Authority in **Managed Detection and Response** Services, protecting the critical data and applications of 1000+ organizations in 70+ countries from known and unknown cyber threats. Founded in 2001, the company's mission is to hunt, investigate and stop cyber threats before they become business disrupting events. Combining cutting-edge machine learning XDR technology, human expertise, and proven security operations leadership, eSentire mitigates business risk, and enables security at scale. The Team eSentire difference means enterprises are protected by the best in the business with a named Cyber Risk Advisor, 24/7 access to SOC Cyber Analysts and Elite Threat Hunters, and industry-leading threat intelligence research from eSentire's Threat Response Unit (TRU). eSentire provides Managed Risk, Managed Detection and Response and Digital Forensic and Incident Response services. For more information, visit www.esentire.com and follow @eSentire.

vmware® Carbon Black

VMware software powers the world's complex digital infrastructure. The company's cloud, app modernization, networking, security, and digital workspace offerings help customers deliver any application on any cloud across any device. Headquartered in Palo Alto, California, VMware is committed to being a force for good, from its breakthrough technology innovations to its global impact.

For more information, please visit <https://www.vmware.com/company.html>

VMware and Carbon Black are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and other jurisdictions.

References

- [1] See Netlogon Elevation of Privilege Vulnerability [Microsoft]
- [2] See Zerologon: Instantly Become Domain Admin by Subverting Netlogon Cryptography (CVE-2020-1472) [Secura]
- [3] The same information, augmented with firsthand research, was also leveraged by eSentire's Threat Response Unit—see Hands-on Threat Research Leads to Resilient Zerologon Detection [eSentire]
- [4] See Privilege Escalation [MITRE]
- [5] See Initial Access [MITRE]
- [6] See Common Vulnerability Scoring System SIG [FIRST]
- [7] See Mitigate Netlogon Elevation of Privilege Vulnerability from August 2020 Patch Tuesday [DHS]
- [8] These figures come from the "Oday 'In the Wild'" spreadsheet maintained by Project Zero
- [9] See HAFNIUM targeting Exchange Servers with 0-day exploits [Microsoft]
- [10] See Analyzing attacks taking advantage of the Exchange Server vulnerabilities [Microsoft]
- [11] Beyond the obvious significance of the brazen theft
- [12] See Hacker Claims to Have Stolen Files Belonging to Prominent Law Firm Jones Day [Wall Street Journal]
- [13] For another example that shows the importance of rapidly operationalizing threat intelligence, see Hands-on Threat Research Leads to Resilient Zerologon Detection [eSentire]