

The Race to the Vulnerable: Measuring the Log4j Shell Incident

Raphael Hiesgen
HAW Hamburg

raphael.hiesgen@haw-hamburg.de

Marcin Nawrocki
Freie Universität Berlin

marcin.nawrocki@fu-berlin.de

Thomas C. Schmidt
HAW Hamburg

t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin

m.waehlich@fu-berlin.de

Abstract—The critical remote-code-execution (RCE) Log4Shell is a severe vulnerability that was disclosed to the public on December 10, 2021. It exploits a bug in the wide-spread Log4j library. Any service that uses the library and exposes an interface to the Internet is potentially vulnerable.

In this paper, we measure the rush of scanners during the two months after the disclosure. We use several vantage points to observe both researchers and attackers. For this purpose, we collect and analyze payloads sent by benign and malicious communication parties, their origins, and churn. We find that the initial rush of scanners quickly ebbed. Especially non-malicious scanners were only interested in the days after the disclosure. In contrast, malicious scanners continue targeting the vulnerability.

Index Terms—Log4j, Log4Shell, Scanning, Security, Network Telescope

I. INTRODUCTION

The world runs on software. With that comes a constant threat of new vulnerabilities which affect widely deployed implementations. In the past, Heartbleed (2014) [1], [2], Rowhammer (2015) [3], and Spectre (2018) [4]–[6] disrupted the operation of systems around the world. The most recent vulnerability in this line is Log4Shell (2021). It enables remote code execution through vulnerable applications by injecting a prepared string into the omnipresent Log4j library.

For hours, days, or weeks after the disclosure of such a vulnerability a race starts between multiple parties. Attackers want to exploit systems before they get patched. Operators and other companies need to update their systems before they get compromised. Security firms monitor malicious activities and researchers aim to understand details and assess the scale of the vulnerability. Releasing a patch and protective measures alongside a public disclosure is a necessary step to enable admins to secure their systems, but no guarantee exists for a timely implementation.

As an integral part of Java applications, Log4jShell can be exploited in a variety of application-dependent ways. We focus our observations on network requests. We collect and analyze data over a two months period using reactive network telescopes at four vantage points on two continents.

After recapitulating the events close to the Log4Shell disclosure in Section II and introducing our data set in Section III, we trace the scanners in Section IV. We find that scanning started on the day of the disclosure and spiked about a week

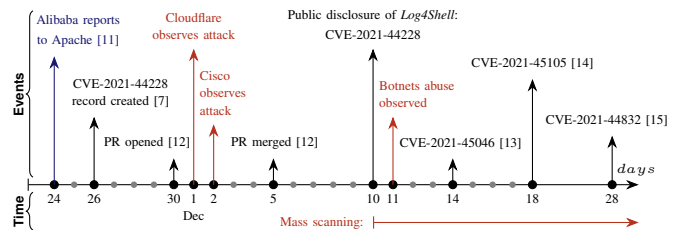


Fig. 1: The unfolding of the Log4Shell vulnerability from reporting to the consequences in 2021.

later with a focus on HTTP related ports. Next, in Section V we examine the collected payloads before focusing on the URLs central to the exploit in Section VI. Finally, Section VII discusses and concludes our findings.

II. BACKGROUND AND RELATED WORK

The Log4Shell vulnerability (CVE-2021-44228 [7]) in the popular Log4j library was publicly disclosed on Dec 10, 2021 by Apache alongside a fix in Log4j library version 2.15.0. The bug allows for remote code execution (RCE) by injecting prepared strings into the logging library. NIST [8] classified this threat as a critical vulnerability with the highest severity rating. An impact assessment by Google [9] estimates that 4% (or 17,000 packages) on the Maven Central were affected, either directly or via dependencies. This is twice the impact of an average packet (mean 2%, median 0.1%), which highlights the popularity of Log4j.

From a high level perspective, the exploit works by injecting a formatted message into the logging component, which then interprets and executes the message. Specifically, the Log4j library supports format messages that are evaluated by the library and can be used to add additional information to log messages, such as the Java version. One of the services that can be used for runtime evaluations is JNDI [10], the Java Naming and Directory Interface, which in turn can query a variety of lookup services. Log4Shell focuses on the services LDAP and RMI for injecting code and infecting the local machine.

A. A Brief History of the Log4Shell Incident

Figure 1 shows the timeline of observations around the Log4Shell incident. The vulnerability was originally reported to Apache on Nov 24 by the Alibaba Cloud Security

Team [11]. The Chinese-based firm quickly faced consequences of reporting the vulnerability directly to Apache [16], [17] instead of contacting national authorities first. A CVE record was created on Nov 26 [7] but not published until public disclosure on Dec 10. In the meantime, a pull request (PR) to address the vulnerability was opened on Nov 30 [12] and merged five days later. Cloudflare notes that they observed the first exploit one day after the PR on Dec 1st [18]. Cisco observed an exploitation attempt a day later on Dec 2nd, as published in their Talos Blog [19]. Both companies report that widespread scanning started on Dec 10.

Within a day of the public release, the Mirai and Muhstik botnets, crypto miners, and other malware were observed to use the exploit for propagation [20], [21]. Microsoft further reports that nation-state attackers experiment with the exploit and integrate it into their activities [22].

The first fix was insufficient and more exploits followed (CVE-2021-45046 [13], CVE-2021-45105 [14], CVE-2021-44832 [15])—none of them as critical, though, as the first.

B. The Log4Shell Attack: How it Works

The Log4Shell exploit builds upon a JNDI injection vulnerability that was presented at BlackHat in 2016 [23]. JNDI enables queries of lookup services such as LDAP, the RMI Registry, or the DNS and can load Java objects returned by a service at runtime. Since the query argument is a URL, the lookup can be performed on local or remote services. Via this functionality, an attacker who controls the query can thus load arbitrary code from a location under his control—and this is where Log4j comes in.

Log4j comprises capabilities to interpret strings for enriching logged messages with additional information. Examples are lookups of the Java version or the hostname. These interpreted strings are escaped by wrapping them: `${prefix:query}`. In addition to harmless operations, Log4j accepts a prefix that triggers JNDI to perform the lookup, in which case the query includes its own scheme to signify the lookup services used for the query. This extends the known JNDI vulnerability by opening an attack vector via logged messages. Applications that log web requests, usernames, or generally user-controlled input are easy targets as a result—provided they fail to sanitize their inputs.

Figure 2 shows a possible exploit scenario initiated by an attacker. The attacker starts with a scan that sends HTTP requests to web servers ①. Here, the exploit string is placed in the user agent field. An information that operators might log to understand what browsers and operating systems they should support to provide good user experience. In our example, the exploit targets an LDAP lookup via JNDI on an LDAP server controlled by the attacker.

The victim application logs the input string using Log4j ②. Log4j in turn finds the escaped string and uses JNDI to perform the LDAP lookup from the remote address ③. The vulnerable Log4j implementation downloads the Java object prepared by the attacker ④ and loads it locally. This Java object contains a way to run shell commands on the local

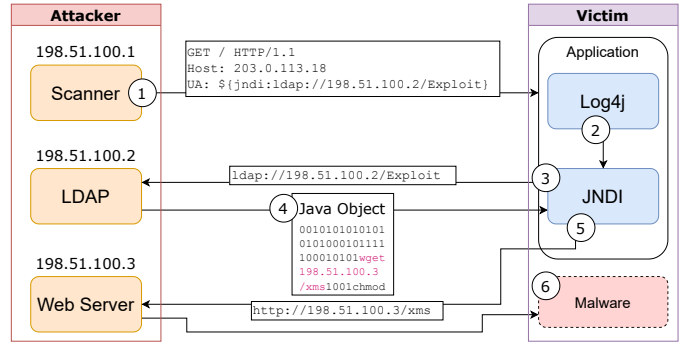


Fig. 2: The Log4Shell exploit tricks the application into loading and executing code from a remote server. It is initially instrumented by an executable log string (here “UA:”).

machine, which downloads the actual malware from a remote server ⑤ and executes it locally ⑥.

Following these steps, an injected payload leads to the execution of arbitrary code prepared by the attacker. The interpretation of strings by Log4j can further be used to obfuscate the payloads [22]. Instead of including the string “ldap” or an address, individual characters are escaped with an operation that leaves them unchanged.

III. OBSERVING LOG4J SCANS: METHODS AND DATA SET

Multiple groups of people hunt for exploitable services: attackers attempting to exploit vulnerable services, researchers who want to examine and analyze, and the security industry, which wants to discover and close vulnerable services before they can be exploited. All of them scan IP space. The time directly after the release of the exploit is the most critical observation period for all, since more and more services are getting patched or taken offline with time. Scanners that want to identify vulnerable services without exploiting them can simply use LDAP servers which do not perform a valid lookup but only log accessing addresses instead. This makes it especially hard to attribute malice based on scans alone.

We observe scan attempts targeting TCP on four /24 IPv4 prefixes. Our vantage points neither host services nor are they part of a larger active network. We deploy Spoki [24], a reactive telescope that interacts with incoming packets in real-time to establish TCP connections, and collect payloads for a few seconds before closing connections. Based on the C++ Actor Framework [25] Spoki is highly scalable and can easily handle thousands of addresses and is deployed at all four vantage points. Our vantage points include one in the US (part of the UCSD Network Telescope [26]) and three in the EU. Two of the vantage points in the EU host neighboring networks (EU VP 2 & VP 3), which are separately announced. We separate the data set by /24 prefixes in our analysis to accommodate for topological differences. We lost data at the US vantage point from Dec 28 to Jan 1.

Payloads targeting the Log4Shell vulnerability contain the escaped JNDI format string. As such, they are detectable by parsing the payloads if a few obfuscation techniques are taken

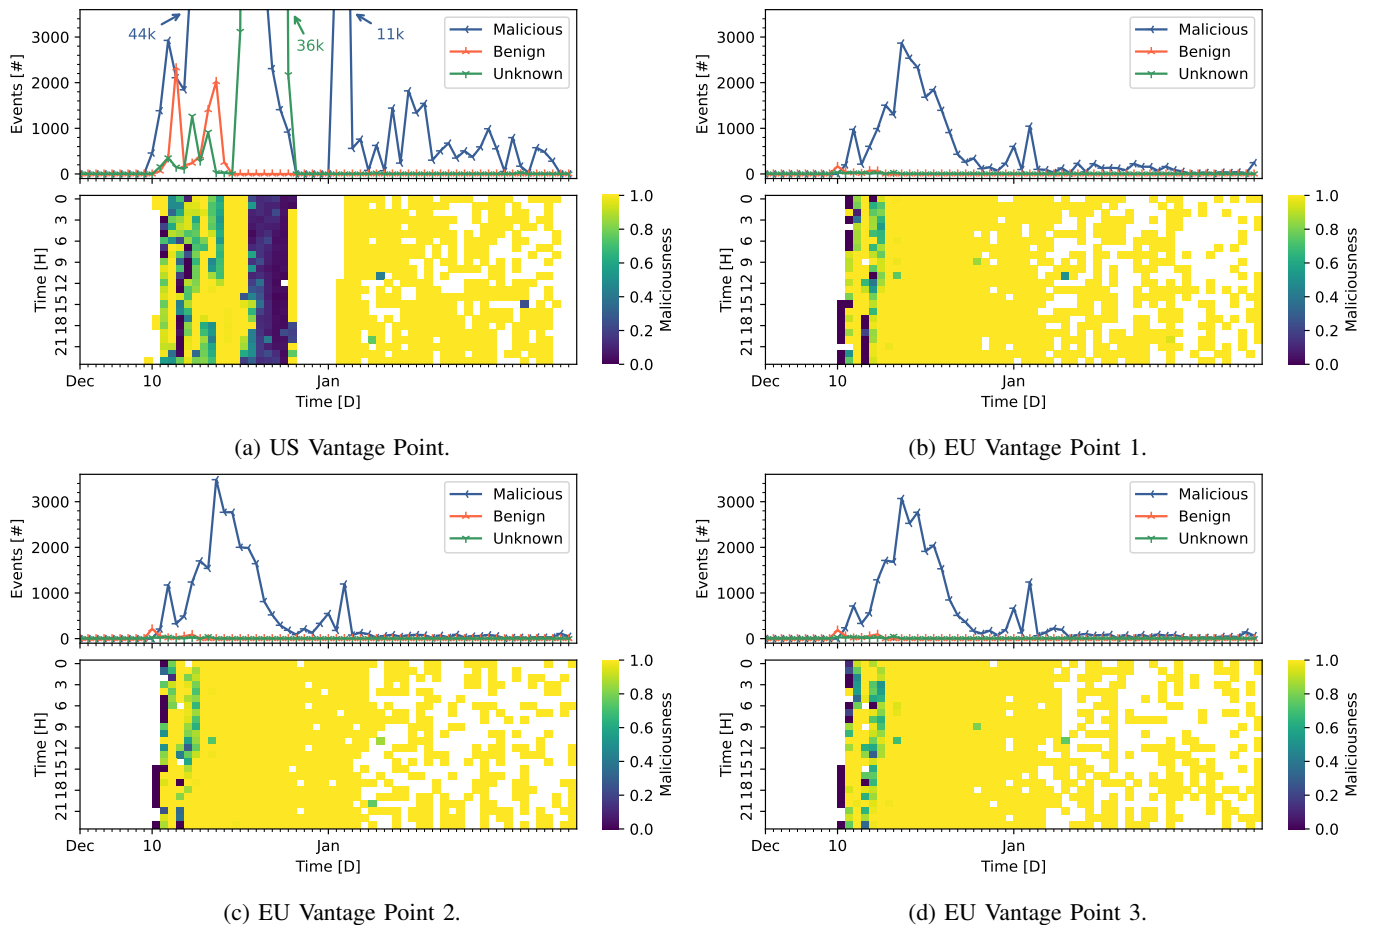


Fig. 3: The intensity and maliciousness of scanners targeting the Log4Shell vulnerability during Dec’21 and Jan’22.

into account. We use an open source tool [27] to filter the payloads we collect. We use MaxMind [28] for geolocating IPs. While geolocation is not highly accurate, mapping to countries is still reasonable [29]. PeeringDB [30] gives us an overview over the types of networks we observe and GreyNoise [31] provides additional threat intelligence information.

IV. SCANNERS

In this section we examine the scanners of Log4Shell scans. These are the sources in step ①, see Figure 2.

A. Overview

On the day of its disclosure (Dec 10, 2021) we observe the first scans targeting the Log4Shell vulnerability at the EU vantage points. In the US first scans started a day earlier on the eve of Dec 9 (UTC). We collect all events, *i.e.*, payloads that include a JNDI exploit string, and classify the sources using GreyNoise [31] into the categories *malicious*, *benign*, and *unknown*. We cannot easily sort *unknown* sources into the other two categories because Log4j payloads sent during step ① do not compromise the system. Still, two payloads that strongly correlate with malicious behavior, see Section VI, were used to further classify sources as malicious.

Figure 3 displays the scan activity over the months Dec’21 and Jan’22. The upper graphs show the time series per source type and day while the heat maps below visualize malicious intensity, calculated from the share of malicious events among all events per hour. The event count in the US graph grows to roughly 44k during the peak of the *malicious* events while the *unknown* events in the subsequent peak first hit roughly 36k.

Mid-December is the period with the highest activity. Noticeably, the only benign events are registered in the first weeks after the disclosure. Towards January the events per day drop to around 100 in the EU, and 2000 or less in the US. The interest in scanning for the vulnerability spiked about a week after the public disclosure. During this time all vantage points receive a strong malicious activity. The US vantage point observes a large share of unknown events as well.

The first scans start on the evening of Dec 9 (UTC+0) at the US vantage point. The EU only sees the first packets nearly a day later (3PM, Dec 10) than the US (11PM, Dec 9). Moreover, the scans in the EU start with scans from BinaryEdge (<https://www.binaryedge.io>), a benign threat-intelligence provider—here from DigitalOcean hosts based in the US. At the US vantage point the first scans originate from one source classified as malicious, in a UK AS.

The heat maps reveal a period of mixed activity before malicious sources take over—with the exception of the unknown sources in the US around mid-December. While the benign actors are fast in the EU, they quickly lose interest. Here, researchers and threat intelligence providers have room to grow: start fast and perform continuous measurements.

From the observed 2023 sources, 1516 were exclusively seen in the US, and 123 only in the three EU telescopes. The remaining 384 addresses were observed at both vantage points. EU VP 1 and 2 bring about ~ 25 exclusive addresses, EU VP 3 only observed addresses also seen at other telescopes.

B. Event Peaks

Shortly after the disclosure all vantage points observe benign scanners. At the US VP we observe two spikes up to 2000 packets on Dec 13 & 18. 90% of these packets are from a single AS, Alpha Strike (<https://www.alphastrike.io>). We observe scans focusing on HTTP-related ports such as 8080, 8081, as well as port 8983.

A very high spike in malicious events was observed in the US on Dec 19 & 20. During these two days, the vantage point records more than 40k packets per day. 80% of them originate from a single IP endpoint in Russia. This is one order of magnitude more than observed from any other source during this time. All these events target port 8080. Two days later, a second spike occurs, this time from unknown sources. During these 5 days the US VP observed more than 30k packets per day. As a result the maliciousness decreases, which is visually striking in the heat map as a dark stripe from Dec 22 to 26.

Once again, we observe a large share of events (76%) originating from a single, unclassified IP address of the same Russian hoster of the previous, malicious address that caused the spike days earlier. In contrast to the previous scan these events are split between TCP ports 8080 (18%) and 5480 (75%). A closer inspection of these two events reveals that the first source scans more aggressively but over a shorter time period. The hourly packet rate is up to twice as high. In contrast, the unknown source sends at a lower rate over a longer time period, adding up to more than twice as many packets in total. Although their payloads have similarities—they exfiltrate the domain, computer name, OS, and java version—they do not show direct intent to perform RCE via Java objects. We also find no specific overlap in the payload strings and do not see sufficient evidence to classify the source as malicious.

A malicious spike in events can be observed at all three EU VPs on Dec 18. These spikes reach about 3k events per day. When grouping by AS, two ASes take more than 10% share: an American hoster and Chinese ISP. Both scan a range of HTTP-related ports as well as ports 5480 and 7547. These ASes show up at the US VP with a slightly higher event count. Their share is comparatively small due to the high traffic volume from the aforementioned two Russian addresses.

Smaller event spikes from malicious sources are seen at the European VPs on Jan 1 & 3. These originate from a single hoster IP address. These probes focus on HTTP-related ports. The much larger spike in the US is caused by the same address.

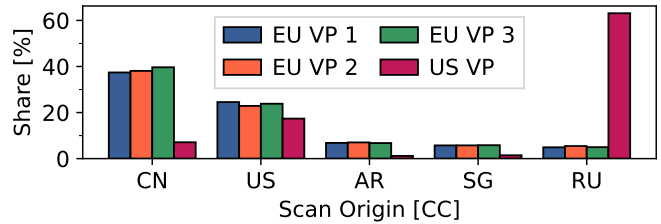


Fig. 4: The top 5 countries with at least 5% share. Russian scanners focus on the US.

The long tail of smaller spikes in the US originates from a few different ASes. Five make up more than 10% of the events each, four of which are hosters. Most scans focus again on HTTP-related ports. Among the lesser scanned ports one stands out: 25565, used for MySQL and Minecraft servers.

C. Who is Scanning?

Geolocation (see Section III) based on the source addresses shows that scans originate to a large extent from three countries. China and the US make up more than 60% of the events at each vantage point in the EU. In line with the traffic spikes caused by Russian sources in the US, Russia originates more than 60% of the events observed in the US and around 5% at the EU vantage points. Figure 4 summarizes the shares of the top five contributing countries.

To determine the type of network, from which scans originate, we query PeeringDB [30]. Table I summarizes the results. 15% to 40% of networks do not have an entry or are labeled as *Not Disclosed*. At all vantage points the largest share falls into the category *Content*, which includes hosting services. Three hosting ASes stand out: Two from the US and one from Russia, which predominantly scans the US VP.

The second-largest share goes to transit and access networks. Traditionally, these are more the home of botnets, which scan from infected hosts [32], [33]. Education and business networks tend to scan because of infected host or with the goal to learn about the new vulnerability.

Educational networks perform scans for research purposes or host compromised machines. The small share of education networks in the US is likely a lack of data in PeeringDB. Although it could further hint that US institutions prefer VMs in the cloud for measurements and infrastructure.

TABLE I: Overview over the network types from PeeringDB. Hosting providers originate the largest identifiable share.

	EU			US
	VP 1	VP2	VP3	VP 1
Hosting	34%	34%	35%	79%
Transit/Access	22%	23%	22%	4%
Education	7%	7%	7%	1%
Enterprise	<1%			<1%
Other	<1%	<1%	<1%	<1%
Unknown	36%	35%	36%	15%

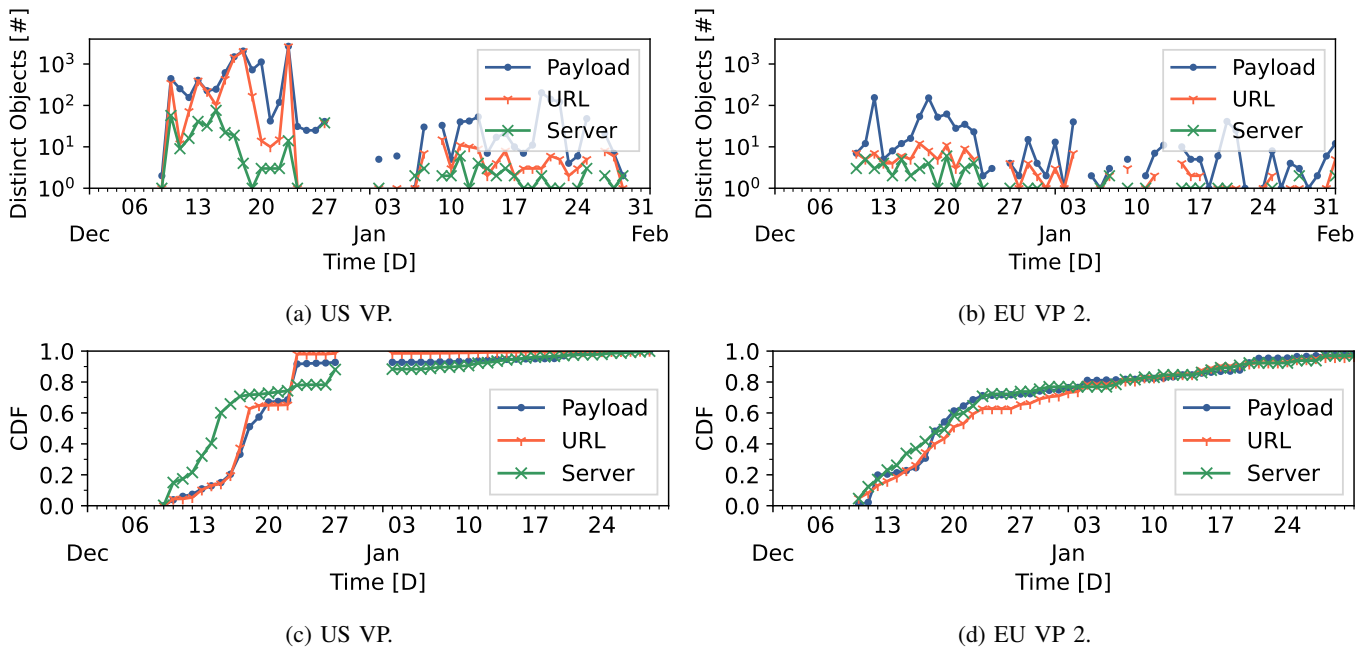


Fig. 5: The relationship between payloads, URLs, and servers over time. (a) and (b) count the distinct objects per day. (c) and (d) aggregate distinct objects over the two-month period in a CDF. EU VPs are represented by EU VP 2 (Dec’21 to Jan’22).

D. What are Scanners Targeting?

The majority of scans aim for HTTP-related ports, such as 80, 8080, 8000. One of the highly-active Russian scanners also focuses on 5480. The three most popular ports account for more than 50%, likewise the top ten ports account for more than 85% of the events at each VP. Overall, we observe between 36 and 48 different ports as targets at each VP.

A port that stands out among the top ten in the EU is 7547. This port is associated with the TR-069 vulnerability related to home routers. While this port is frequently scanned (about 4%), Java is generally not the best fit for home routers and thus Log4j is not likely to be used in such an environment. Presumably, scanners are just testing the port in passing.

Scan packets do not uniformly distribute across the addresses of our prefixes. At every vantage point the difference between the least scanned address and the most scanned address is at least a factor of 1.5 (US) and 2.5 (EU).

V. PAYLOADS OF THE SCANNERS

All observed Log4Shell events contain a payload within the JNDI lookup string. Hence, the number of payloads we collected equals the number of events described in Section IV. The basic idea of such payloads with embedded exploits was explained in Section II-B. We now analyze the collected payloads, *i.e.*, the data scanners send in step ①, in detail to learn about what they impact, how payloads change over time, which protocols scanners use, and the JNDI URL placement.

A. Temporal Development

Section IV-A revealed how scans reached a quick peak before they rapidly decline and settle in a low volume. We cannot

tell from the event counts alone whether scanners keep their setups and continue to run unaltered, *i.e.*, reusing payloads, or whether they purposefully rotate payloads. The variation in payloads over time can help to answer this question.

Figure 5 depicts the evolution of distinct payloads, URLs, and hosting servers over our observation period. Subfigures 5a and 5b show the distinct counts per day. In contrast, Subfigures 5c and 5d show the CDF of distinct objects over the complete period. The payloads are the data we receive directly via TCP, URLs are the escaped JNDI URLs in the payloads, and servers are solely the addresses in those URLs, *i.e.*, the host and port information of LDAP or RMI server controlled by the attacker. Because payloads and URLs often include the address of the victim, *e.g.*, in the host parameter of the HTTP header, we replace the IP addresses from our subnets with a static string. This way, payloads that only differ by their destination will coincide.

The event spikes in Figures 5a and 5b correlate with the spikes about a week after the disclosure, see Figure 3. Note that the graph has a logarithmic y-axis. At both vantage points the distinct payloads count is much higher than the server count. Scanners use different payloads for the same injection string—likely to test the behavior of the scanned service as it might differ by application which data is logged. Similarly, the number of distinct URLs is higher than the number of distinct server addresses. In those cases, attackers re-use the same server, but with different paths. In practice this means that individual servers are used to distribute multiple types of downloads or malware. On some days the number of payloads equals the URL count. Here, scanners encoded additional data in the URL, such as a HEX-encoded JSON payload that

TABLE II: The most common header locations to store the JNDI URL. (UA: *User-Agent*, Auth: *Authorization*, X-API-Ver: *X-API-Version*)

	EU						US	
	VP 1		VP2		VP3		VP 1	
1. UA	23%	UA	22%	Auth.	23%	UA	11%	
2. Auth.	20%	Auth.	21%	UA	22%	Path	9%	
3. Path	13%	Path	14%	Path	14%	Cookie	6%	
4. Cookie	10%	Cookie	11%	Cookie	11%	Auth.	6%	
5. X-API-Ver.	10%	X-API-Ver.	8%	X-API-Ver.	9%	Referer	6%	

includes the destination and other information.

We now analyze the evolution of the attack ecosystem over the complete time frame, *i.e.*, the new payloads, URLs, and servers. Figures 5c and 5d plot the corresponding CDFs. Noticeably, the behavior of distinct payloads and infrastructure largely differ between the US and the EU vantage points: The US graph payloads and URLs has more extreme steps, one from Dec 17 to 20 and one from Dec 22 to Dec 23. These jumps lead up to the high spikes in malicious and later in unknown events. In contrast, the server line rises much earlier. The rise in payloads and URLs does coincide. This indicates varying attacks in the US backed by the same infrastructure—supposedly, attackers adapted their payloads and URLs to explore the attack surface.

At the EU VP, all three measures occur in line. There are two steeper increases in payloads shortly after mass scanning started on the Dec 12 and again on Dec 18. Still, we do not see a divergence as in the US. Payloads, URLs, and servers behave similarly, rising on similar days. Scanners vary their payloads. Most frequently we observe classic horizontal scans and scanners sending different payloads to a single endpoint before moving on.

B. Exploit String Placement

While focusing on HTTP related ports we observe that payloads of the scanners are either GET (between 91% and 98%) or PUT requests (remaining). Note that we cannot observe all application-specific payloads as they may be encoded in an application-specific format. Nevertheless, binary protocols that contain a matching string in ASCII would still be detected.

Exploit strings in payloads become visible at various locations. This reflects attackers who still search for the best attack method, as well as attackers that try to alter their payloads continuously to evade detection. The most popular placements in HTTP headers are summarized in Table II.

Regularly logged locations such as the *User-Agent* (UA) and *Authentication* (Auth) header fields are popular targets. The UA, which can be customized for specific purposes [34], identifies the client. Websites depend on it to deliver matching content. As a result, it is often logged to keep statistics on users. Authentication information are similarly important for debugging and access control. Since HTTP header manipulation is a known issue [35], logging any header fields should be done with care.

The lower shares in the US stem from the high variety in header configurations and fields sent by the Russian scanners. While we saw around 45 different header fields in the EU, the US VP observed 157 distinct fields. During January the field count shrinks to 13, favoring User Agent and X-API-Version.

The two Russian scanners follow different scanning styles. The one IP address tagged as malicious appears to partly randomize the headers. All payloads are GET requests that include the user agent `Go-http-client`, a `Host` field, the “randomized” field with the JNDI string, and end with `Connection: close`. Exception are payloads that contain the JNDI string as the user agent, in which case only the `Host` field follows. Not only do they change their fields but choose from several simple obfuscated URLs.

In contrast, the other Russian scanner rotates payloads systematically. It sends the same payload to a number of addresses before changing a single characteristic, such as the header field that contains the JNDI string or the obfuscation method. It looks like a nested loop that rotates the obfuscation in the inner loop and the header field in the outer one. The user agent is only included when it contains the JNDI string.

The obfuscation approaches we observe are built on the substitution used for the vulnerability itself. There are a few easy obfuscations that hide the strings “jndi”, “ldap” or other parts of the URL, hiding keywords that can otherwise be found by simple pattern matching, thus making the payloads harder to detect. As an example, the string `lower:j` will substitute to “j”. As the obfuscation is very easy to achieve and the general JNDI exploit is not new, obfuscated payloads quickly appeared in our data. At the EU vantage points, the first obfuscated payloads showed up shortly after midnight on Dec 11. In the US it took a day longer.

VI. EXAMINING THE JNDI/LDAP EXPLOITATION

In this section, we investigate the final steps ③ to ⑤ of the attack, namely the malware requests triggered via JNDI. We also acquire and inspect the malware distributed via Log4Shell. The URLs used to query via JNDI have four parts: a scheme, a host, a port, and a path.

A. Analyzing the URLs

Schemes. Tricking the victim into downloading Java objects via JNDI is central to the exploit, see step ③. JNDI supports a range of services, but only a few are used for Log4Shell.

Figure 6 shows the event count for each scheme we observed (note the logarithmic y-axis). The EU vantage points see LDAP almost exclusively. RMI occurs a handful of times at both vantage points. Aside from the events on Jan 16, which originate from the same ASes, there is no correlation between the VPs. As a third scheme, EU VP 3 observed a single HTTP request. The website behind the address gives security advice on how to patch Log4j (Feb’22). The name in the reverse DNS record of the scanning node matches BinaryEdge. In the US we observe two spikes of DNS schemes in January. 95% of these originate from the Alphas Strike AS, matching the benign spikes in Figure 1.

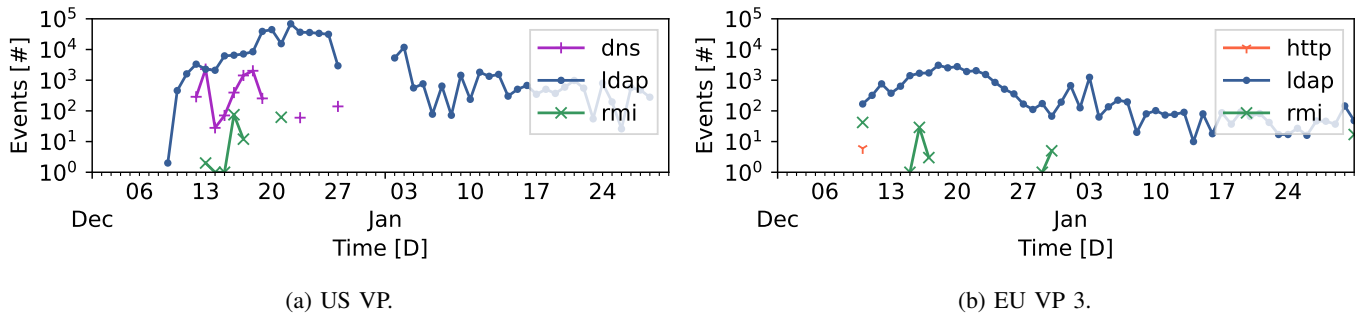


Fig. 6: Schemes in JNDI URLs over time. LDAP is used nearly exclusively (Dec’21 to Jan’22).

Aside from LDAP, attacks over RMI received attention in the media [36]. We only observed three (US) request with the URL mentioned in the article. These attacks may have been focused and did not hit our VPs in bulk.

Hosts. PeeringDB lacks information about nearly half of the hosting networks we observe. We notice that corresponding servers are mostly located in two ASes, both from Estonia (one provider shows up with two ASes, one of which is Ukrainian). Manual search identified both as hosters. Relabeling them accordingly, the servers we observe are mostly located in what PeeringDB labels as *Content* ASes, *i.e.*, hosters (EU: 70%, US: 80%). The large share of hosters fits the distribution model for malware. These servers need to be reachable from everywhere with a high uptime to allow compromised machines to download malware. Transit and access networks together make up roughly 20% at the EU VPs and 5% at the US VP.

The most popular location for servers are displayed in Figure 7. While Estonia runs top in the EU, Russia takes the biggest chunk for the US by far. Notably, none of the Russian servers were observed in our EU vantage points.

Next we compare the location of the scanners to the server locations. In the EU region the most common combination is scanning from China and hosting in Estonia (about 15%). Second rank is using the US for both scanning and hosting (roughly 12% of the events) These combinations are dwarfed by the Russian scanners observed in the US (70%), which use the same IP for scanning and hosting the LDAP server. Here, the combinations US and US (5%), China and Estonia (3%) follow on places two and three.

Ports. Less than two percent of the LDAP servers used for

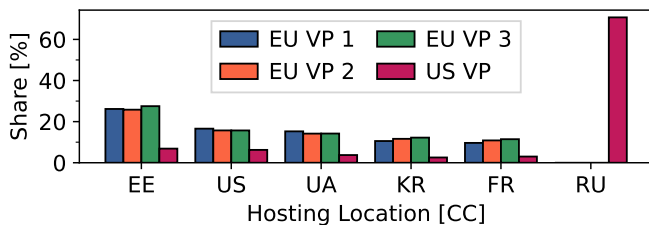


Fig. 7: The top 5 countries that host Log4Shell servers. Russia only takes a share in the US.

the attacks are bound to the default LDAP port (389). Instead, the port most commonly used is 1389. The share ranges from 93% to 96% at all vantage points. Other ports in use are 2420 and 80 in the EU, each between 1% and 2%. In the US we observe 12344 at around 2% as well. Such a high share could hint at common tools or tutorials used by the attackers.

Paths. Except for a single path, observed paths do not conform to the RFC that defines them [37] as they do not include a valid distinguished name. Two paths stand out among the LDAP URLs. First is `/Exploit`, which makes up the largest share at all vantage points with 70% to 80% in the EU and 20% in the US. A second group of paths shares the segments `Command/Base64` followed by a base64-encoded segment. This group takes the second-largest share. These paths begin in a variety of ways, potentially hinting at their purpose, such as `TomcatBypass` or `GroovyBypass`. Decoding the base64 segment reveals script code (mostly bash) that downloads an executable via HTTP to run locally.

At first glance, encoding commands in a URL path is an odd choice. However, the second piece of the puzzle is an LDAP server implementation, which dynamically builds Java objects that will run the base64-encoded command. These servers can be found on GitHub in several repositories, although the original¹ is no longer available. This server—aptly named `JNDIExploit`—binds local ports for LDAP and HTTP. It responds to a variety of queries that include the `Base64` fragment and match the paths we observed. The most common variant we see is the `TomcatBypass`.

Having such tooling at hand makes it easy for attackers to set up the attack and run it. We confirmed that servers used in attacks exhibit this dynamic behavior by sending a custom base64-encoded command. We received the Java object in response of our custom code. The `JNDIExploit` repositories use port 1389 as the default for LDAP, which matches our observation of common LDAP ports in the URLs.

B. Downloading Malware

The Java objects are just an intermediate step to the goal of the attack: infecting the host with malware. We follow this path to find out which malware is distributed this way. When running the downloader in Feb’22, most servers were

¹<https://github.com/feihong-cs/JNDIExploit>

no longer available. We successfully download nine distinct Java objects, compare step ④.

The LDAP answers contain two important keys: `javaClassName` and `javaSerializedData`. In all cases the class name is set to `java.lang.String`. The objects we collected match the objects built by the JNDIExploit LDAP server, see Section VI-A. One was built for the “Groovy bypass” and the remaining were built for the “Tomcat bypass”. The serialized objects look like a Java `StringRefAddr` object.

Both bypasses use slightly different ways to run code. The Tomcat bypass instantiates a script engine to run JavaScript code while the Groovy bypass builds on Groovy itself. The script code is encoded in the serialized objects as ASCII. One of the Tomcat samples executes PowerShell code, which likely targets Windows. While most of the other scripts include a mechanism to determine the local OS, e.g., by checking the direction of slashes in a file path, they execute bash commands either way, and are thus unlikely to run on Windows.

We extract the download commands in our small samples set by hand and download three different binaries, compare step ⑤. (Four connections failed and three of the objects contain the same scripts, although they differ slightly.) The hashes of all downloads are registered on VirusTotal [38], where they were first submitted between the mid and end of Jan’22, i.e., while Log4Shell attacks were taking place. Two of the samples are scripts while one is a 32-bit ELF binary.

The malware acquired via the PowerShell code is itself a PowerShell script and downloads a binary with the name of a known crypto miner. Although the other shell script looks more sophisticated—it stops local programs, downloads an uninstaller to remove software, adds new cron tab entries, and tries to make its way into connected hosts via ssh—its goal is similar: installing a crypto miner.

Understanding the base64-encoded commands opens another way to acquire malware. Instead of taking the round-trip via the LDAP server, which embeds them into a Java object, the commands can be decoded directly. Via this method we can acquire an additional binary: a 64-bit executable ELF file first registered on VirusTotal at the beginning of Dec’21. The low yield matches the churn in URLs we observe in Section V.

C. Locating Malicious LDAP Servers

We finally explore the malicious server infrastructures by active scans. To this end, we utilize the fact that the Log4Shell exploit requires publicly reachable servers that return Java objects, compare step ③ and ④ in Figure 2.

Malicious servers predominantly listen on port 1389 coupled with a non-standard LDAP behavior. This is why we use ZMap [39] to scan TCP/1389 for open ports. We then identify unsecured LDAP servers by performing LDAP bind operations, which should fail on servers enforcing authentication. In a next step we query for the two most common paths observed: `/Exploit` and a path with `Base64` string.

We find 5.1M servers responding to SYNs, but only 1,110 allow unauthorized LDAP-binding. 81 servers return answers

for `/Exploit`, and 179 for `Base64`. These sets overlap, which leads to 183 malicious LDAP servers in total (16%). Comparing to Figure 5, we infer that the number of daily used and dormant servers differ by an order of magnitude.

We collect six Java objects via the `/Exploit` path and 97 objects via the `base64` path. Their general structure matches the objects from Section VI-B. Since `base64` requests are expected to encode our command they should not return any malicious objects. Six objects are an exception. Here, servers return the same payload for both paths, likely a static response. We identify one downloader for a 64 Bit ELF binary, an object that runs PowerShell code, and one broken object. The remaining three payloads do not include script code.

VII. DISCUSSION AND CONCLUSION

The ease with which Log4Shell can turn into exploits is part of what make it as critical as it is. Access to some public interface suffices to abuse it. YouTube videos explain details and give guidance on how to apply the tools. The underlying problem is input sanitization—a common challenge that has plagued the industry in the form of SQL injections for years. User-supplied input must be treated with caution.

On a positive note, the vulnerability saw wide coverage online as blog posts, lists of vulnerable applications, and detection tools were quickly published. At the same time official organizations published reports and issued warnings.

In this paper, we observed Log4Shell scanning through reactive network telescopes immediately after the disclosure. Most noticeable were the large spikes of malicious events about a week after. These hit all vantage points but particularly targeted the US, giving the scans a geographical focus. More vantage points are necessary to confirm this as a global pattern.

Our analysis showed common characteristics such as two prevalent path fragments: `Exploit` and `base64`-encoded commands. Such similarities hint at shared tools or tutorials. This assumption is further strengthened by a common LDAP port we observed during the majority of attacks and active scans, which aids us in identifying LDAP servers used for attacks.

Although the overall scan rate decreased significantly during our observation period, malicious scanners continue to scan for the vulnerability and transport new payloads.

The methodology presented in this paper is limited to the observation of scans. As such, we cannot measure the success rate of attackers, but observing scanning behavior can be an expressive indicator for the liveliness of the scene. A quick decrease in scanning activity may also reflect that vulnerable services are thinning out.

While the list of affected software contains many popular applications [40], the long term effects are yet unclear. Many applications quickly saw patches, but their rollout will eventually slow down. Due to the many attack vectors it is hard to find all active vulnerable applications using a single methodology. The combination of complementary approaches will be part of our future work.

ACKNOWLEDGMENTS

We thank CAIDA for providing access to the UCSD Network Telescope and the network operators who support us deploying our telescope in Europe. We appreciate fruitful discussions with Johannes Klick and Stephan Lau. This work was partly supported by the German Federal Ministry of Education and Research (BMBF) within the project *PRIME*net.

REFERENCES

- [1] Z. Durumeric, F. Li, J. Kasten, J. Amann, J. Beekman, M. Payer, N. Weaver, D. Adrian, V. Paxson, M. Bailey, and J. A. Halderman, “The Matter of Heartbleed,” in *Proc. of ACM IMC*. New York, NY, USA: ACM, 2014, pp. 475–488.
- [2] The MITRE Corporation, “CVE-2014-0160,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2014-0160>.
- [3] The MITRE Corporation, “CVE-2015-0565,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-0565>.
- [4] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom, “Spectre Attacks: Exploiting Speculative Execution,” *Commun. ACM*, vol. 63, no. 7, pp. 93–101, Jun 2020.
- [5] The MITRE Corporation, “CVE-2017-5753,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2017-5753>.
- [6] The MITRE Corporation, “CVE-2017-5715,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=cve-2017-5715>.
- [7] The MITRE Corporation, “CVE-2021-44228,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-44228>.
- [8] NIST, “CVE-2021-44228 Detail,” Feb 2022, <https://nvd.nist.gov/vuln/detail/CVE-2021-44228>.
- [9] Google, “Understanding the Impact of Apache Log4j Vulnerability,” Feb 2022, <https://security.googleblog.com/2021/12/understanding-impact-of-apache-log4j.html>.
- [10] Apache, “JNDI Lookup plugin support,” Feb 2022, <https://issues.apache.org/jira/browse/LOG4J2-313>.
- [11] Apache, “Apache Log4j Security Vulnerabilities,” Feb 2022, <https://logging.apache.org/log4j/2.x/security.html>.
- [12] R. Goers, “Restrict LDAP access via JNDI,” Feb 2022, <https://github.com/apache/logging-log4j2/pull/608>.
- [13] The MITRE Corporation, “CVE-2021-45046,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-45046>.
- [14] The MITRE Corporation, “CVE-2021-45105,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-45105>.
- [15] The MITRE Corporation, “CVE-2021-44832,” Feb 2022, <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-44832>.
- [16] South China Morning Post, “Apache Log4j bug: China’s industry ministry pulls support from Alibaba Cloud for not reporting flaw to government first,” Feb 2022, <https://www.scmp.com/tech/big-tech/article/3160670/apache-log4j-bug-chinas-industry-ministry-pulls-support-alibaba-cloud>.
- [17] ZDNet, “Chinese regulators suspend Alibaba Cloud over failure to report Log4j vulnerability,” Feb 2022, <https://www.zdnet.com/article/log4j-chinese-regulators-suspend-alibaba-partnership-over-failure-to-report-vulnerability/>.
- [18] Matthew Prince (Cloudflare), Feb 2022, <https://twitter.com/eastdakota/status/1469800951351427073>.
- [19] Cisco Talos, “Threat Advisory: Critical Apache Log4j vulnerability being exploited in the wild,” Feb 2022, <https://blog.talosintelligence.com/2021/12/apache-log4j-rce-vulnerability.html>.
- [20] 360 Netlab, “Threat Alert: Log4j Vulnerability Has Been adopted by two Linux Botnets,” Feb 2022, <https://blog.netlab.360.com/threat-alert-log4j-vulnerability-has-been-adopted-by-two-linux-botnets/>.
- [21] Juniper, “Log4j Attack Payloads In The Wild,” Feb 2022, <https://blogs.juniper.net/en-us/security/in-the-wild-log4j-attack-payloads>.
- [22] Microsoft, “Guidance for preventing, detecting, and hunting for exploitation of the Log4j 2 vulnerability,” Feb 2022, <https://www.microsoft.com/security/blog/2021/12/11/guidance-for-preventing-detecting-and-hunting-for-cve-2021-44228-log4j-2-exploitation/>.
- [23] A. Muñoz and O. Mirosh, “A Journey From JNDI/LDAP Manipulation to Remote Code Execution Dream Land,” Feb 2022, <https://www.blackhat.com/docs/us-16/materials/us-16-Munoz-A-Journey-From-JNDI-LDAP-Manipulation-To-RCE-wp.pdf>.
- [24] R. Hiesgen, M. Nawrocki, A. King, A. Dainotti, T. C. Schmidt, and M. Wählisch, “Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope,” in *Proc. of 31st USENIX Security Symposium*. Berkeley, CA, USA: USENIX Association, August 2022.
- [25] D. Charousset, R. Hiesgen, and T. C. Schmidt, “CAF - The C++ Actor Framework for Scalable and Resource-efficient Applications,” in *Proc. of the 5th ACM SIGPLAN Conf. on Systems, Programming, and Applications (SPLASH ’14), Workshop AGERE!* New York, NY, USA: ACM, Oct. 2014, pp. 15–28.
- [26] CAIDA, “The UCSD Network Telescope,” Website, 2012, last Access: May 2022. [Online]. Available: https://www.caida.org/projects/network_telescope/
- [27] F. Roth, “log4shell-detector,” Dec 2021, <https://github.com/Neo23x0/log4shell-detector>.
- [28] “MaxMind - GeoLite Country.” [Online]. Available: http://www.maxmind.com/app/geoip_country
- [29] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, “IP Geolocation Databases: Unreliable?” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 2, pp. 53–56, Apr 2011.
- [30] PeeringDB, “The Interconnection Database,” <https://www.peeringdb.com/>, 2019.
- [31] GreyNoise, “IP Lookup API,” <https://greynoise.io>.
- [32] A. Dainotti, A. King, K. Claffy, F. Papale, and A. Pescapè, “Analysis of a ‘/0’ Stealth Scan from a Botnet,” *IEEE/ACM Transactions on Networking*, vol. 23, no. 2, pp. 341–354, Apr 2015.
- [33] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, “Understanding the Mirai Botnet,” in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug 2017, pp. 1093–1110.
- [34] Google, “Google-Crawler (User-Agents),” May 2022, <https://developers.google.com/search/docs/advanced/crawling/overview-google-crawlers>.
- [35] G. Tyson, S. Huang, F. Cuadrado, I. Castro, V. C. Perta, A. Sathiseelan, and S. Uhlig, “Exploring HTTP Header Manipulation In-The-Wild,” in *Proc. of the WWW*. Republic and Canton of Geneva, CHE: ACM, 2017, pp. 451–458.
- [36] Juniper, “Log4j Vulnerability: Attackers Shift Focus From LDAP to RMI,” Feb 2022, <https://blogs.juniper.net/en-us/threat-research/log4j-vulnerability-attackers-shift-focus-from-ldap-to-rmi>.
- [37] M. Smith and T. Howes, “Lightweight Directory Access Protocol (LDAP): Uniform Resource Locator,” IETF, RFC 4516, June 2006.
- [38] “VirusTotal,” <https://www.virustotal.com/>.
- [39] Z. Durumeric, E. Wustrow, and J. A. Halderman, “ZMap: Fast Internet-wide scanning and its security applications,” in *Proc. of the 22nd USENIX Security Symposium*. USENIX Assoc., Aug. 2013.
- [40] National Cyber Security Centrum, “Log4shell vulnerabilities,” Feb 2022, <https://github.com/NCSC-NL/log4shell>.