

Delays have Dangerous Ends: Slow HTTP/2 DoS attacks into the Wild and their Real-Time Detection using Event Sequence Analysis

Nikhil Tripathi, *Member, IEEE*

Abstract—The robustness principle, written by Jon Postel in an early version of TCP implementation, states that the communicating entities should be *liberal* while accepting the data. Several entities on the Internet do follow this principle. For instance, in this work, we show that many popular web servers on the Internet are generous as they wait for a substantial time period to receive the remaining portion of an incomplete web request. Unfortunately, this behavior also makes them vulnerable to a class of cyber attacks, commonly known as Slow Rate DoS attacks. HTTP/2, the recent version of HTTP, is recently found vulnerable to these attacks. However, the impact of Slow HTTP/2 DoS attacks on real web servers on the Internet has not been studied yet. Also, to the best of our knowledge, there is no defense scheme known to detect Slow Rate DoS attacks against HTTP/2 in real-time. To bridge these gaps, we first test the behavior of HTTP/2 supporting web servers on the Internet against Slow HTTP/2 DoS attacks. Subsequently, we propose a scheme to detect these attacks in real-time. We show that the proposed detection scheme can detect attacks in real-time with high accuracy and marginal computational overhead.

Index Terms—HTTP/2, Slow Rate DoS attacks, Anomaly detection, DoS

1 INTRODUCTION

APPLICATION layer Denial-of-Service (DoS) attacks are a newer class of DoS attacks that exploit vulnerabilities present in the application layer protocols such as HTTP, NTP, DHCP [1], [2], [3]. These attacks target specific services (e.g., web, mail, file transfer) running on a victim computer with a relatively slighter impact on the network resources [4], [5]. One famous example of application layer DoS attacks is the Slow Rate DoS attack. Slow Rate DoS attack requires sending multiple incomplete requests to the victim server to consume its connection queue space. Once this finite space is consumed, the server can not entertain further incoming requests. A few instances of Slow Rate DoS attacks have also been encountered in the past [6], [7].

HTTP/1.1 is one of the most studied protocols against Slow Rate DoS attacks [1], [7]. HTTP/2, the successor of HTTP/1.1, was launched in May 2015 [8], and our HTTP/2 usage measurement study (see Section 4) shows that approximately 71% of the websites serve web content using HTTP/2. The operational differences between HTTP/1.1 and HTTP/2 make the Slow Rate DoS attacks against HTTP/1.1 [1] ineffective against HTTP/2. Likewise, the defense mechanisms known to counter Slow Rate DoS attacks against HTTP/1.1 are ineffective in detecting Slow Rate DoS attacks against HTTP/2. In one of our prior works [9], we performed the vulnerability assessment of HTTP/2, and proposed five new Slow Rate DoS attacks that are effective against it. However, we tested the attacks in a controlled lab setup only, and we did not observe the behavior of real HTTP/2 web servers against them. Thus, it is essential to evaluate the effects of Slow HTTP/2 DoS

attacks on the real web servers deployed on the Internet. Taking motivation from it, in this work, we test the behavior of HTTP/2 supporting web servers on the Internet against Slow HTTP/2 DoS attacks, and showed that several popular web servers are vulnerable to them¹. We also propose a detection mechanism in this work to proactively detect these attacks. We make the following specific contributions in this paper:

- We perform a comprehensive Internet-wide study by scanning the top 500K Alexa websites to find how many of them support HTTP/2. Our study shows that approximately 71% of the top 500K websites support HTTP/2.
- We test the behavior of HTTP/2 supporting web servers on the Internet against different Slow HTTP/2 DoS attacks and show that several of them are vulnerable to the attacks.
- We propose an event sequence analysis-based detection scheme to detect Slow HTTP/2 DoS attacks in real-time.
- We test the detection performance of the proposed scheme in a real network and show that it can detect attacks with very high accuracy and marginal computational overhead.

The rest of the paper is organized as follows. We present the HTTP/2 overview and the literature review in Sections 2 and 3. In Sections 4 and 5, we then discuss our study to measure HTTP/2 usage on the Internet and the empirical evaluation of Slow HTTP/2 DoS attacks, respectively. In Section 6, we discuss the proposed detection scheme and

• N. Tripathi is with the Computer Science & Engineering Group at Indian Institute of Information Technology, Sri City, Chittoor, India 517646. E-mail: nikhil.t@iits.in

1. The empirical evaluation results are also presented in our recent preliminary work [10].

present the experiments conducted to test the detection performance of the proposed scheme in Section 7. Finally, the paper is concluded in Section 8.

2 PRELIMINARIES

This section presents some background knowledge on HTTP/2 and Slow Rate DoS attacks.

2.1 HTTP/2

HTTP/2 is defined in RFC 7540 [8] and recently became a standard in May 2015. HTTP/2 operation differs significantly from its predecessor HTTP/1.1 [11], [12]. HTTP/2 was proposed keeping in mind the inability of HTTP/1.1 to utilize the TCP's transmission capacity efficiently. HTTP/2 achieves higher efficiency by using several virtual connections known as *streams* over a single TCP connection. HTTP/2 also mitigates head-of-line blocking of the web requests that persists in HTTP/1.1. HTTP/2 also implements its own flow control mechanism at the application layer. It is required to prevent the communicating entities from overwhelming each other by sending data simultaneously through multiple streams over a single TCP connection. HTTP/2 also involves using different types of frames for different purposes [8]. For example, HTTP/2 HEADERS and CONTINUATION frames are used to transmit the header part of the web request, while HTTP/2 DATA frames are used to transmit the message body (e.g., HTML code of a web page). HTTP/2 SETTINGS frames are used to negotiate connection parameters, and HTTP/2 GOAWAY frames are used to close the connection. HTTP/2 WINDOW_UPDATE frame indicates the size of data (in bytes) that the sender can accept from the receiver in addition to the current window size. These frames are carried inside the TCP payload and should not be confused with the link-layer frames. Interested readers are requested to refer to RFC 7540 [8] for the list of HTTP/2 frames and their usage.

2.2 Slow Rate DoS Attacks

Slow Rate DoS attacks [13] require sending multiple specially crafted incomplete requests to a victim server. On receiving such requests, the server stores them in its connection queue space and waits for a time duration \mathcal{T} (predefined in the configuration file) in the hope of receiving the remaining portions of the web request. However, the attacker never sends the remaining portion of the complete requests. To consume the server's connection queue space, the attacker establishes enough connections, and from each connection, it sends incomplete requests. Once the connection queue space is consumed, the server can not entertain further incoming requests, resulting in DoS scenario. It is not a good approach to mitigate these attacks by terminating the TCP connections with a connection duration greater than a predefined threshold. This is because incomplete requests are not uncommon on the Internet [1], [9], and such an approach may also terminate the connections from the benign clients with poor TCP connection quality. Examples of such clients are remote clients with high latency and clients on low-grade cellular or satellite networks.

3 RELATED WORKS

Several works in the literature discuss DoS/DDoS attacks against HTTP/1.1 and the defense mechanisms to counter those attacks [7], [13], [14]. However, the same does not hold true for its successor, HTTP/2. Since HTTP/2 is a recently standardized protocol, only a few works discuss possible attacks against HTTP/2 and the defense mechanisms to counter them [7]. Moreover, since HTTP/2 operation differs significantly from HTTP/1.1 [11], [12], the existing techniques to detect Slow Rate DoS attacks against HTTP/1.1 [1], [15], [16], [17], [18] are ineffective in detecting Slow Rate DoS attacks against HTTP/2. As a result, researchers in the security community recently proposed a few defense approaches to detect Slow HTTP/2 DoS attacks. This section first mentions DoS/DDoS attacks known against HTTP/2 and, subsequently, the defense mechanisms to counter them.

3.1 DoS/DDoS attacks against HTTP/2

Imperva [19] disclosed some implementation vulnerabilities in HTTP/2, which can be exploited for creating situations such as Blue Screen of Death and arbitrary code execution at the client side. However, these vulnerabilities existed in the protocol implementations but not in the protocol itself. Post the vulnerability exposure; the affected implementations were patched to mitigate these vulnerabilities. Beckett and Sezer [11], [12] studied how HTTP/2 functionalities can be exploited to launch flood and amplification-based DDoS attacks. Adi et al. [20] tested how the HTTP/2 web servers behave if control frame flooding is launched against them. Praseed and Thilagam [21], [22] demonstrated that the request multiplexing feature of HTTP/2 can be exploited to launch asymmetric DDoS attacks on the web servers. In particular, the attacker sends high workload web requests resulting in heavy computational overhead at the web server. The multiplexed asymmetric attack may result in either partial DoS, also known as Reduction-of-Quality (RoQ), or a complete DoS scenario. The attacks discussed in these works are distributed in fashion and, thus, typically fall under the Distributed Denial-of-Service (DDoS) category [7]. Since these attacks require sending a large number of web requests to the server to render it unusable, they are relatively difficult to launch and easier to detect [23].

Unlike traditional DDoS attacks, Slow HTTP/2 DoS attacks [9] need small computational power for their execution at the attacker's side. They involve sending very few requests to the web server to prevent it from serving benign clients [7], [23]. Since these attacks generate a smaller amount of traffic, they are stealthier and, thus, difficult to detect as compared to traditional DDoS attacks [7], [13], [23]. **Taking motivation from this, we restrict our focus only to Slow HTTP/2 DoS attacks in this paper.**

Based on the HTTP/2 frame parameters that are altered, there can be different variants of Slow HTTP/2 DoS attacks. Interested readers are requested to refer to our previous work [9] for a detailed working of these attacks. For completeness, we present a summary of these attacks below:

3.1.1 Advertising Zero Window Size (Attack-1)

In this attack, the attacker sends an HTTP/2 SETTINGS frame to the server to inform that it does not have any window space available to receive the data. The server does not

send any data back to the attacker on receiving this frame unless it receives an HTTP/2 WINDOW_UPDATE frame from the attacker. However, the attacker never sends the required WINDOW_UPDATE frame to the server forcing the latter to wait.

3.1.2 Incomplete POST Request Message Body (Attack-2)

HTTP/2 uses DATA frames to carry the message body of a web request (e.g., values in an HTML form). A DATA frame defines an END_STREAM flag. This flag, when *set*, indicates that the frame carries the entire message body and will not be followed by any other DATA frame. However, to launch this attack, the attacker sends a DATA frame having END_STREAM flag *reset*. On receiving this frame, the server waits to receive more DATA frames such that the last DATA frame must have the END_STREAM flag set. However, the attacker never sends the subsequent DATA frames, forcing the web server to wait for the complete message body.

3.1.3 Sending Connection Preface Only (Attack-3)

Connection preface is used to inform a web server that HTTP/2 will be used for further communication. The attacker only sends the connection preface to launch this attack and does not send any subsequent HTTP/2 frames. On receiving the connection preface only, the server waits to receive further HTTP/2 frames belonging to the web request.

3.1.4 Incomplete GET/POST Request Header (Attack-4)

HTTP/2 HEADERS frames carry the header part of the web requests. If the header is large enough not to fit in a single HEADERS frame, it is split into smaller parts such that the first header part is sent into the HEADERS frame, while the remaining parts are sent into the CONTINUATION frames. A HEADERS frame defines an END_HEADERS flag. This flag, when *set*, indicates that the HEADERS frame carries the entire header, and any CONTINUATION frame will not follow it. However, to launch this attack, the attacker sends a HEADERS frame having the END_HEADERS flag *reset*. On receiving this frame, the server waits to receive CONTINUATION frames such that the last CONTINUATION frame must have the END_HEADERS flag set. However, the attacker never sends the subsequent CONTINUATION frames, forcing the web server to wait for the complete header.

3.1.5 Unacknowledged SETTINGS frame (Attack-5)

On receiving an HTTP/2 SETTINGS frame, a recipient (client/server) must acknowledge it by sending an empty SETTINGS frame back to the sender. The attacker does not acknowledge an HTTP/2 SETTINGS frame sent by the server to launch this attack. As a result, the web server keeps waiting to receive the SETTINGS frame acknowledgment.

3.2 Known Defense Mechanisms

Several works in the literature discuss defense strategies to counter attacks against HTTP/1.1. However, those strategies can not detect attacks against HTTP/2 due to the several operational differences between HTTP/1.1 and HTTP/2 [11], [12]. Moreover, since HTTP/2 is a newly standardized

TABLE 1: Comparison with prior works

References	Empirical Evaluation of Slow HTTP/2 DoS	Detection of Slow HTTP/2 DoS
Adi et al. [20]	✗	✗
Tripathi and Hubballi [9]	✗	✓(only offline)
Praseed and Thilagam [21]	✗	✗
Praseed and Thilagam [22]	✗	✗
Beckett and Sezer [11]	✗	✗
Beckett and Sezer [12]	✗	✗
Praseed and Thilagam [24]	✗	✗
Tripathi and Shaji [10]	✓	✗
This paper	✓	✓

protocol, only a few existing works focused on HTTP/2 security. In one of our previous works [9], we proposed an offline anomaly detection scheme to detect if a web server received incomplete web requests in a particular time interval. However, the detection scheme could detect if the interval contains attack traffic only after it is elapsed. Thus, it could not detect the attacks in real-time. In another work, Praseed and Thilagam [24] proposed an approach to model an HTTP/2 request set as a fuzzy multiset and subsequently assign a trust score to distinguish between valid and invalid request sets. This approach could detect the multiplexed asymmetric DDoS attack with decent accuracy. However, it was not designed to detect Slow HTTP/2 DoS attacks.

3.3 Empirical Evaluation of Slow HTTP/2 DoS Attacks

In our recent preliminary work [10], we tested the behavior of HTTP/2 supporting web servers on the Internet against Slow HTTP/2 DoS attacks. We showed that several popular web servers are vulnerable to these attacks. However, we did not present any defense approach in that work which can be deployed to counter Slow HTTP/2 DoS attacks.

Table 1 briefly summarizes how our work differs from the prior works related to HTTP/2 security.

4 MEASURING HTTP/2 USAGE ON THE INTERNET

In September 2021, we conducted a comprehensive study to measure HTTP/2 usage on the Internet. We took into account the top 500,000 Alexa domains in this study and tested if they support HTTP/2. For this purpose, we wrote a python script that communicates with the web server of each domain in the Alexa list. In particular, the python script sends a ‘Client Hello’ message to the web server as part of the standard TLS handshake. Through this message, the script informs the web server about the application protocols that it supports. Examples of these application protocols are HTTP/2 over TLS (known as h2), HTTP/2 over cleartext (known as h2c), and HTTP/1.1. In response, the server also sends a ‘Server Hello’ message to the script, thereby informing about the application protocols that it supports.

We conducted this measurement study into four iterations. In the first iteration, we collected HTTP/2 usage statistics by performing TLS handshakes with the web

servers of the top 500K domains in the Alexa list. In this iteration, we received responses from 445315 web servers, out of which 335854 web servers were found to support HTTP/2. Moreover, 109461 web servers were found to support HTTP/1.1. The remaining 54685 web servers did not respond in the first iteration for two reasons. First, the TCP connections were timed out as there was no response from some web servers. Second, the authoritative name servers for some of the domains have an A record for the hostname `www` only. For example, the authoritative name servers for a domain, say, `iiits.ac.in` have the record `www.iiits.ac.in A 103.21.58.130`, but they do not have the record `iiits.ac.in A 103.21.58.130`. Thus, the web server `103.21.58.130` is reachable only when a web request is sent to the Fully Qualified Domain Name (FQDN) `www.iiits.ac.in`. We also observed that many domains support HTTP/1.1, but they serve web content over HTTP/2 if a user visits the FQDN `www.<domain>` as the web server mapped to the FQDN `www.<domain>` supports HTTP/2, while the web server mapped to `<domain>` supports HTTP/1.1 only. The result of the first iteration measurement is shown in Table 2.

In the second iteration, we appended the `www` keyword to all those domains (=164146) whose web servers either supported HTTP/1.1 (=109461) or did not respond (=54685) in the first iteration. This provided us FQDNs of the format `www.<domain>`. Subsequently, we collected HTTP/2 usage statistics by performing TLS handshakes with the web servers corresponding to these FQDNs. The result of the second iteration measurement is shown in Table 2. It can be noticed from the table that 17220 more HTTP/2 supporting domains were found in the second iteration. Also, the number of domains that either supported HTTP/1.1 (=101841) or were unresponsive (=45085) was decreased in the second iteration.

The popular web clients such as Chrome, Firefox, and Safari support h2 only; however, the debugging and benchmarking tools such as curl and h2load support h2 as well as h2c. Thus, popular server-side implementations such as Apache simultaneously support both h2c and h2 configurations to avoid compatibility issues. In the first and second iterations, we measured the usage of h2 on the Internet using the HTTPS scheme. In the third iteration, we measured the usage of h2c on the Internet using the HTTP scheme. We checked the HTTP/2 support only for those domains which either supported HTTP/1.1 or did not respond in the second iteration. To measure the h2c usage statistics, our script sends a cleartext HTTP/2 web request to a web server after the 3-way TCP handshake and subsequently checks if the response from the web server is an HTTP/2 response. The result of the third iteration measurement is shown in Table 2. It can be noticed from the table that no response was received for 144135 domains. Thus, in the fourth iteration, we appended the `www` keyword to the unresponsive domains in the third iteration and obtained FQDNs of the format `www.<domain>`. Subsequently, we collected HTTP/2 statistics by sending cleartext HTTP/2 web requests to the web servers corresponding to these FQDNs. The result of the fourth iteration measurement is shown in Table 2. It can be noticed from the table that the web servers corresponding to **356417** (sum of the numbers

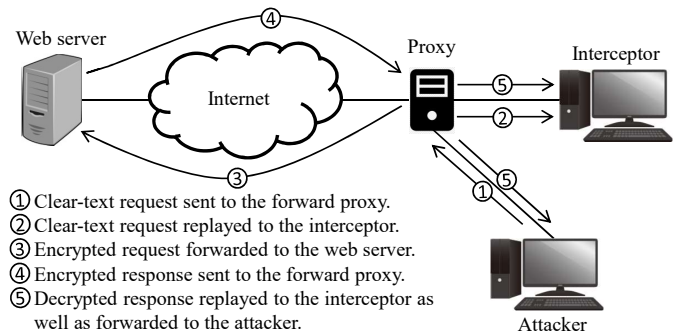


Fig. 1: Testbed setup

shown in bold fonts in Table 2) domains were found to support HTTP/2.

5 EMPIRICAL EVALUATION OF THE ATTACKS

This section presents an empirical evaluation of Slow HTTP/2 DoS attacks. We created a testbed setup for this evaluation, as shown in Figure 1.

5.1 Testbed Setup

Our testbed had four entities - 1) an attacker machine, 2) a transparent forward proxy, 3) an interceptor, and 4) the target web server.

Attacker: We designated a computer as the attacker machine and executed on it the attack scripts written in Python. These attack scripts sent the web servers exactly one incomplete HTTP/2 request of each attack (Attack-1 to 5) type. We sent *only* one incomplete request to observe the servers' behavior, and we did not intend to launch the attacks against the servers by sending multiple attack requests.

Transparent Forward Proxy: As discussed in the previous section, most of the web servers on the Internet support HTTP/2 over TLS (h2). Thus, to intercept and monitor otherwise encrypted HTTPS traffic, we configured a transparent forward proxy [25] in front of the attacker machine such that the proxy connects to the web servers on behalf of the attacker. The proxy takes cleartext HTTP/2 requests from the attacker as input and establishes TLS connections² with the web server of the domain present in the `HOST` field of the HTTP/2 requests. Subsequently, it forwards the encrypted web request to the server. On receiving the encrypted HTTP/2 response from the server, the proxy decrypts the response and forwards it to the attacker. It also replays the cleartext HTTP/2 traffic from the web server and the attacker to the interceptor, as shown in Figure 1. Before replaying the traffic to the interceptor, the proxy modifies the IP addresses appropriately such that the ultimate endpoints of the traffic become the web server and the attacker. The proxy uses the `PCAP-OVER-IP` method to replay the traffic to the interceptor on a designated port 57012. In this way, we intercepted the cleartext HTTP/2 traffic so that it could be analyzed later.

Interceptor: We designated a computer as the interceptor

2. The TLS connections are established only with h2 servers. In the case of h2c servers, the proxy simply forwarded the cleartext requests.

TABLE 2: HTTP usage statistics. h2: HTTP/2 over TLS, h2c: HTTP/2 over cleartext, http/1.1: HTTP/1.1, NA: Not Applicable

		1st iteration (https://<domain>)	2nd iteration (https://www.<domain>)	3rd iteration (http://<domain>)	4th iteration (http://www.<domain>)
Responded	h2	335854	017220	NA	NA
	h2c	NA	NA	002791	000552
	http/1.1	109461	101841	000000	000000
Not responded		054685	045085	144135	143583
Total		500000	164146	146926	144135




Fig. 2: CDF plot of the waiting time of top 500K Alexa web-sites. Please refer to the color version for better illustration.

that intercepted the cleartext HTTP/2 traffic replayed by the proxy. The interceptor machine captured the inbound traffic to port 57012 using tcpdump and stored it in a pcap file.

Target Web Server: The target web server belongs to the 356417 HTTP/2 supporting web servers that we found on the Internet.

5.2 HTTP/2 Traffic Analysis

The captured HTTP/2 traffic at the interceptor was analyzed to compute the connection waiting time for different target web servers. For each flow³ in the pcap file, we calculated the time difference between the connection establishment (3-way handshake) and its termination (FIN-ACK packet exchange). In this way, we obtained the time the web server waited before closing the connection from which an incomplete request was sent. A CDF plot of the duration for which different web servers wait in case of each attack is shown in Figure 2. We can notice from the figure that approximately 3% of the web servers waited for more than 360 seconds before closing a connection from which an attack-3 type web request was sent. However, approximately 50% of the web servers were found to wait for the same duration before closing the connection from which an attack-1 type web request was sent. Since attack-1 can make more web servers

3. We differentiated the TCP flows using four parameters - source IP, source port, destination IP, and destination port.

wait for a particular time duration, it is more effective than attack-3. The other similar inferences can also be drawn from the graph shown in Figure 2.

6 PROPOSED DETECTION SCHEME

An HTTP/2 request involves exchanging different types of HTTP/2 frames, as discussed in Section 2.1. If we treat the exchange of these frames as unique events that occur one after the other, we can obtain an event sequence that is possible during the interaction between a client and an HTTP/2 server. Since the frames exchanged during the HTTP/2 request depend upon various parameters such as the user’s browsing behavior and available window size at the receiver, several distinct event sequences are possible. All possible normal event sequences can be stored into a database of characteristic normal patterns (observed sequences of exchanged HTTP/2 frames) and subsequently be used to detect anomalous sequences. An HTTP/2 request is considered as attack request if the event sequence generated corresponding to it does not exist in the normal database and vice versa. However, simply checking the non-existence of an event sequence in the database to detect attack requests may result in several false positives. This is because there can be many possible unique event sequences (see Section 7.1.1), and maintaining a database of all the sequences to capture the complete normal behavior of HTTP/2 would be a tedious task.

To avoid this issue, we maintain a database of possible lookahead pairs extracted from the possible event sequences. Later in Section 7.1.1, we show that the database of possible **lookahead pairs** can capture the normal behavior of HTTP/2 with more completeness than the database of possible **event sequences**. Once we obtain the normal database of lookahead pairs during the learning phase, we assign a mismatch score to an event sequence during the detection phase based on the number of lookahead pairs extracted from it that do not exist in the normal database. The event sequence is declared anomalous (and thus, the request from which it is generated, is declared as attack request) only when its mismatch score exceeds a predefined threshold value.

6.1 Learning Phase

During the learning phase, our detection scheme builds a normal database of lookahead pairs extracted from the possible event sequences and a database of maximum normal delay possible between the occurrences of two consecutive events. Algorithm 1 describes the procedure for generating

Algorithm 1 Learning Phase

Input: \mathcal{F} - Clear-text HTTP/2 flows; n - Window size
Output: $\mathcal{D}_{lookahead}$, \mathcal{D}_{delay}

```

1:  $\mathcal{D}_{lookahead}=\{\}; \mathcal{D}_{delay}=\{\}; i=1;$ 
2: for  $flow$  in  $\mathcal{F}$  do
3:    $seq='Start \rightarrow *';$ 
4:   for  $frame$  in  $flow$  do
5:      $e=frame.translateToEvent();$ 
6:      $seq+=' \rightarrow e \rightarrow *';$ 
7:   end for
8:    $seq+=' \rightarrow End';$ 
9:    $l=seq.extractLookaheadPairs();$ 
10:   $\mathcal{D}_{lookahead}.append(l)$ 
11:   $events = seq.extractEvents();$ 
12:  while  $i \leq len(events) - 1$  do
13:     $delay=events[i].time()-events[i-1].time();$ 
14:     $str='events[i-1] \rightarrow events[i]';$ 
15:    if  $str$  not in  $\mathcal{D}_{delay}$  then  $\mathcal{D}_{delay}[str]=delay;$ 
16:    else
17:       $\mathcal{D}_{delay}[str]=max(delay, \mathcal{D}_{delay}[str]);$ 
18:    end if
19:     $i+=1;$ 
20:  end while
21: end for
22: return  $(\mathcal{D}_{lookahead}, \mathcal{D}_{delay});$ 

```

these databases during the learning phase. Below we discuss the procedure for generating these databases in detail.

6.1.1 Bulding the lookahead pairs database ($\mathcal{D}_{lookahead}$)

During the learning phase, we first collect normal HTTP/2 traffic and subsequently construct individual TCP flows by combining TCP payloads of all the packets exchanged after the connection establishment and before the connection termination. Once flows are constructed, they are parsed using Deep Packet Inspection (DPI) to spot HTTP/2 frames. These frames are then translated into events, as shown in Table 3. After translating the HTTP/2 frames into the corresponding events, we append the events one after the another to generate an event sequence. For the event sequence generation, we consider a TCP connection's 3-way handshake as the beginning of the flow (represented by event *Start*) and the connection termination (represented by event *End*) as the end of the flow. One sample event sequence for a TCP flow $IP_A:port_A \leftrightarrow IP_B:port_B$ carrying an HTTP/2 request is as follows:

$$Start \rightarrow * \rightarrow Pref \rightarrow * \rightarrow Max_Con_Strm!0 \rightarrow Ini_Win_Size!0 \rightarrow * \rightarrow win_size_incr!0$$

In this sequence, '*' event corresponds to the end of an HTTP2 frame, while ' \rightarrow ' is not an event and only represents a transition from one event to another.

Once an event sequence is generated, the lookahead pairs are extracted from it and stored in $\mathcal{D}_{lookahead}$. For this, we slide a window of size $n + 1$ across the event sequence and capture which events follow which within the sliding window. For example, if we consider $n = 3$ and the following event sequence:

TABLE 3: HTTP/2 frames and translated events

Frame	Event
3-way handshake	<i>Start</i>
Connection preface	<i>Pref</i>
DATA frame with End Stream (ES) reset	<i>Data_frame!ES</i>
DATA frame with ES set	<i>Data_frame_ES</i>
HEADERS frame with ES and End Header (EH) reset	<i>Hdr_frame!(ESEH)</i>
HEADERS frame with ES reset and EH set	<i>Hdr_frame!ES_EH</i>
HEADERS frame with ES and End Header (EH) set	<i>Hdr_frame_(ESEH)</i>
Acknowledgement to SETTINGS frame	<i>Settings_ACK</i>
SETTINGS frame with non-zero length	<i>Settings_UNACK</i>
SETTINGS frame having initial window size=0	<i>Ini_Win_Size = 0</i>
SETTINGS frame having initial window size!0	<i>Ini_Win_Size!0</i>
SETTINGS frame having maximum concurrent stream!0	<i>Max_Con_Strm!0</i>
SETTINGS frame having maximum concurrent stream=0	<i>Max_Con_Strm = 0</i>
WINDOW frame with increment in window size!0	<i>win_size_incr!0</i>
WINDOW frame with increment in window size=0	<i>win_size_incr = 0</i>
GOAWAY frame	<i>GOAWAY</i>
CONTINUATION frame	<i>CONTINUATION</i>
Connection termination	<i>End</i>

TABLE 4: Lookahead pairs for w_1

	Index-1	-2	-3
<i>Start</i>	*	<i>Max_Con_Strm!0</i>	<i>Ini_Win_Size!0</i>
*	<i>Max_Con_Strm!0</i>	<i>Ini_Win_Size!0</i>	
<i>Max_Con_Strm!0</i>	<i>Ini_Win_Size!0</i>		

$$Start \rightarrow * \rightarrow Max_Con_Strm!0 \rightarrow Ini_Win_Size!0 \rightarrow * \rightarrow win_size_incr!0 \rightarrow * \rightarrow Pref$$

For each event, we capture the event that comes after it at spot 1, spot 2, and so on, up to spot n , as we slide the window across the above event sequence. For the first window w_1 , from index 1 (*) to index 3 (*Ini_Win_Size!0*) in the sequence, we obtain the database having six lookahead pairs as shown in Table 4.

Similarly, after sliding the window across the complete event sequence, we obtain the lookahead pairs as shown in Table 5. A lookahead pair is stored in the database in the form of ' $event[i] : event[j], k$ ' where $k \leq n$ (window size) and corresponds to the distance between the i^{th} and j^{th} events in a sequence. Once the database of lookahead pairs ($\mathcal{D}_{lookahead}$) is built, our detection scheme refers to it to assign anomaly scores to the event sequences generated from incoming HTTP/2 requests during detection phase.

TABLE 5: Lookahead pairs extracted from the complete event sequence

	Index-1	-2	-3
<i>Start</i>	*	<i>Max_Con_Strm!</i> 0	<i>Ini_Win_Size!</i> 0
*	<i>Max_Con_Strm!</i> 0 ; <i>win_size_incr!</i> 0 ; <i>Pref</i>	<i>Ini_Win_Size!</i> 0 ;*	* ; <i>Pref</i>
<i>Max_Con_Strm!</i> 0	<i>Ini_Win_Size!</i> 0	*	<i>win_size_incr!</i> 0
<i>Ini_Win_Size!</i> 0	*	<i>win_size_incr!</i> 0	*

6.1.2 Bulding the maximum delay database (\mathcal{D}_{delay})

The proposed detection scheme also maintains a database, \mathcal{D}_{delay} , of the maximum normal delay possible between the occurrences of two consecutive events, \mathcal{A} and \mathcal{B} . This is essential to detect the determined attackers who send the benign HTTP/2 frames to the server but with a prolonged delay between them to launch the Slow Rate DoS attacks. An entry in the maximum delay database is stored in the form ' $event[i] \rightarrow event[i+1] = d$ ' where d is the maximum delay between the i^{th} and $(i+1)^{th}$ events in a sequence.

6.2 Detection Phase

Algorithm 2 describes the procedure for detecting an anomalous event sequence during the detection phase. We use different threads for executing the functions responsible for sniffing the incoming packets and extracting events from them, checking timeouts between the events, and assigning anomaly scores to the sequences. However, for brevity, we omit these details in Algorithm 2. The proposed scheme generates an event sequence seq for each incoming HTTP/2 request during the detection phase using the method discussed in Section 6.1.1. During this phase, the scheme also inserts a 'timeout string' after an event in seq if a time duration t passes after the occurrence of such that t is greater than the maximum delay possible between the occurrence of and any other event that can occur immediately after. The maximum possible delay value is retrieved from \mathcal{D}_{delay} database generated during the learning phase, as discussed in Section 6.1.2.

Once the length of seq exceeds the window size n , we start extracting lookahead pairs from seq . Subsequently, we test for the presence or absence of the extracted lookahead pairs in $\mathcal{D}_{lookahead}$ to assign the mismatch score to seq .

6.2.1 Assigning Mismatch Scores

Consider an example where the database contains the entries as shown in Table 5, and we generate the following event sequence from an HTTP/2 request at a particular time instance t_i :

$Start \rightarrow * \rightarrow Max_Con_Strm!0 \rightarrow * \rightarrow$
 $Ini_Win_Size!0 \rightarrow * \rightarrow win_size_incr!0 \rightarrow * \rightarrow$
 $Pref$

For $n=3$, the lookahead pairs extracted from this event sequence are shown in Table 6. Since the lookahead pairs shown in red in Table 6 do not exist in Table 5, these are

Algorithm 2 Detection Phase

Input: $\mathcal{D}_{lookahead}$; \mathcal{D}_{delay} ; n - Window size; - Mismatch score maximum threshold
Input: \mathcal{F} - Incoming TCP flow
Output: D - Declare sequence as anomalous or normal

- 1: $seq = ''$;
- 2: **while** True **do**
- 3: Continuously sniff packets belonging to \mathcal{F} ;
- 4: **for** $frame$ in \mathcal{F} **do**
- 5: $e = frame.translateToEvent()$;
- 6: $seq += e \rightarrow *'$;
- 7: **end for**
- 8: $events = seq.extractEvents()$;
- 9: $last_event = events[-1]$;
- 10: $max_delay = \max(\mathcal{D}_{delay}['last_event \rightarrow \mathcal{E}'])$ where \mathcal{E} belongs to set of all possible events;
- 11: **if** $(current.time() - last_event.time()) > max_delay$ **then** $seq += ' \rightarrow TO_i \rightarrow *'$ where $i = 1, 2, \dots$;
- 12: **end if**
- 13: $mismatch = 0$;
- 14: **if** $len(events) > n$ **then**
- 15: $l = seq.extractLookaheadPairs()$;
- 16: **for** $pair$ in l **do**
- 17: **if** $pair$ not in $\mathcal{D}_{lookahead}$ **then** $mismatch += 1$;
- 18: **end if**
- 19: **end for**
- 20: $mismatch = mismatch / (n * (len(events) - (n + 1) / 2))$;
- 21: **if** $mismatch >$ **then** seq is anomalous;
- 22: **else if** $mismatch <$ and $events[-1] == 'End'$ **then**
- 23: seq is normal;
- 24: **end if**
- 25: **end if**
- 26: **end while**

TABLE 6: Lookahead pairs extracted from the test sequence

	Index-1	-2	-3
<i>Start</i>	*	<i>Max_Con_Strm!</i> 0	*
*	<i>Max_Con_Strm!</i> 0 ; <i>Ini_Win_Size!</i> 0 ; <i>win_size_incr!</i> 0 ; <i>Pref</i>	*	<i>Ini_Win_Size!</i> 0 ; <i>win_size_incr!</i> 0 ; <i>Pref</i>
<i>Max_Con_Strm!</i> 0	*	<i>Ini_Win_Size!</i> 0	*
<i>Ini_Win_Size!</i> 0	*	<i>win_size_incr!</i> 0	*
<i>win_size_incr!</i> 0	*	<i>Pref</i>	

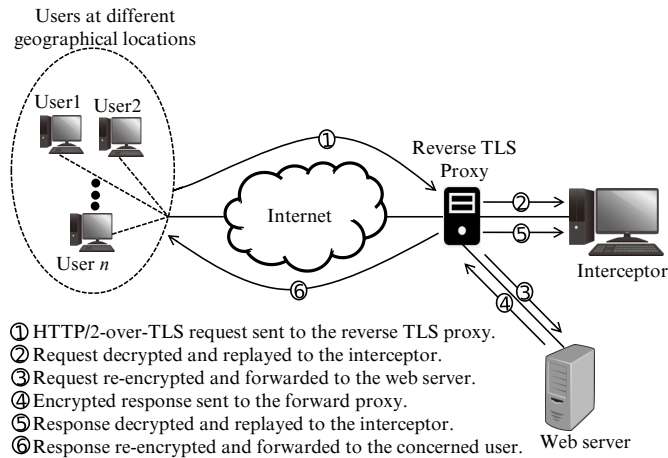


Fig. 3: Testbed setup for the experiments

counted as mismatches. Subsequently, the mismatch score of the event sequence is computed as the ratio of the mismatch counts to the total possible mismatch counts. For instance, if we consider an event sequence of length \mathcal{L} and a lookahead (window size) of n , the maximum pairwise mismatch count is given by

$$n(\mathcal{L}-n)+(n-1)+(n-2)+\dots+1 = n(\mathcal{L}-(n+1)/2)$$

Thus, in the example discussed earlier, the mismatch score of the sequence is $7/21 (= 0.33)$. The mismatch score of the event sequence is computed after every new event appended to the event sequence. It will be continued unless either of the following two conditions is met:

1. The mismatch score of the sequence exceeds the predefined threshold. In this case, the corresponding HTTP/2 request from which the sequence is being generated will be treated as attack request.
2. The *End* event is appended to the event sequence. After appending the *End* event, if the mismatch score of the sequence is less than the predefined threshold, the corresponding HTTP/2 request will be treated as benign. However, if the score exceeds the threshold, the request will be treated as attack request.

7 EXPERIMENTS

7.1 Learning Phase

There are no publicly available dataset for clear-text HTTP/2 traffic (h2c) [24]. Thus, we created a testbed similar to the one shown in Figure 3 to capture clear-text normal HTTP/2 traffic during the learning phase and subsequently used it to build the lookahead pairs, and maximum delay databases:

Web Server: The web server was running Ubuntu 20.04 LTS operating system and configured with Apache 2.4.51 to handle web requests. It had Intel Xeon E3 processor and 64 GB of physical memory. The web server was one of the academic servers of our institution, and it was hosting quiz assignments. The students were asked to attempt the assignments online for their academic grading. During the COVID-19 pandemic, the students visited the website from separate locations, and thus, the collected HTTP/2 traffic exhibited the property of real Internet behavior. We collected

approximately 5 GB HTTP/2 traffic between October 8th, 2021, to October 14th, 2021.

User: The users in our testbed were the students of our institute who accessed the web server to complete their assignments. They were asked to use the latest HTTP/2 supporting web browsers. We also practiced a few non-technical formalities to ensure that the students could not send any malicious traffic to poison the collected traffic.

Reverse TLS Proxy: We configured a reverse TLS proxy application called PolarProxy [25] in front of the web server. This proxy was owned by the website administration and was responsible for the tasks shown in Figure 3.

Interceptor: We designated a computer as the interceptor that intercepted the cleartext HTTP/2 traffic replayed by the proxy. The interceptor machine captured the inbound traffic to port 57012 using tcpdump.

7.1.1 Capturing the normal behavior of HTTP/2


It is essential to decide what size of the database should be built to capture the normal behavior of HTTP/2. For this purpose, we collected 14094 HTTP/2 flows, and for each flow, we generated an event sequence. It resulted in 7221 unique event sequences. We also extracted the lookahead pairs from the unique event sequences and obtained 227, 377, 471, 621, and 737 unique lookahead pairs for window size (n)=3, 4, 5, 6, and 7, respectively. The plots in Figures 4a and 4b show how the number of unique event sequences and lookahead pairs increases as we increase the number of HTTP/2 flows. We can notice from Figure 4a that the number of new event sequences kept increasing with the increase in the number of flows. Thus, it is not easy to estimate how much HTTP/2 traffic should be collected beforehand to obtain the unique event sequences possible in the case of HTTP/2 interactions.

On the contrary, as can be noticed from Figure 4b, virtually no new lookahead pairs were witnessed beyond a point (approximately 9400 flows). Thus, we could assume that the database of lookahead pairs captured the virtually complete normal behavior of HTTP/2 at . From this study, we can also establish that a database of lookahead pairs could capture the normal behavior of HTTP/2 with more completeness than a database of unique event sequences.

7.2 Detection Phase

During the detection phase, we extended the setup shown in Figure 3 by including one more entity and designating it as the attacker. The attacker machine was connected to the Internet and was responsible for sending the web server a small number of attack (Attack-1 to 5) requests at regular intervals. Moreover, the attacker terminated a connection⁴ from which an attack request was sent as soon as 100 seconds progressed after the connection establishment. This was to ensure that such attack requests did not deplete the server's connection queue space. Closing these connections deliberately after 100 seconds did not alter the event sequences generated from the attack requests because our detection scheme could classify such requests as attack or benign way before the aforementioned 100 seconds (see

4. A connection was terminated using FIN-ACK TCP packet.



(a) Number of possible unique event sequences



(b) Number of possible lookahead pairs for different values of n

Fig. 4: Capturing normal HTTP/2 behavior. Please refer to the color version for better illustration.

Section 7.2.3). The benign requests were also generated during the detection phase while the students attempted their assignments.

The reverse TLS proxy in our testbed captured the benign and attack HTTP/2 requests generated by the students and the attacker, respectively, and forwarded their clear-text versions to the interceptor. Our detection scheme running on the interceptor then generated event sequences from the requests and subsequently classified them as normal or anomalous depending on their mismatch scores. We conducted this experiment into five different time intervals such that the window size n was fixed, but the mismatch score threshold was varied across the intervals. In this way, we could analyze the effect of varying n on the detection performance of the scheme. We used four metrics to evaluate the detection performance - Accuracy, False Positive Rate (FPR), Recall, and Precision. The detection performance of the scheme in terms of these four metrics is shown in Table

Fig. 5: CDF plot of the mismatch scores of normal and anomalous sequences for $n=5$. Please refer to the color version for better illustration.

7. We can notice from the table that *Accuracy* and *Precision* slightly increased as we increased n to 0.02 from 0.01. It was because of the reduced number of *FP* cases (a normal sequence incorrectly classified as an anomalous sequence) and the increased number of *TN* cases. Due to the same reason, *FPR* decreased as we increased n to 0.02 from 0.01. However, as we further increased n , *Accuracy* decreased because of more *FN* cases (an anomalous sequence incorrectly classified as a normal sequence). Moreover, *Recall* and *Accuracy* decreased drastically as we increased n to 0.03 from 0.02. It was because the anomalous sequences with a mismatch score in the range $[0.02, 0.03]$ were considered as normal resulting in a large number of *FN* cases. Due to these misclassifications, we obtained lower detection accuracy. The same holds true as we further increased the value of n .

7.2.1 Mismatch scores of normal and anomalous event sequences

It can be noticed from Figure 5 that the number of anomalous sequences with a mismatch score in the range $[0, 0.02]$ is almost the same across this range, and the mismatch scores of virtually all normal sequences laid in the range $[0, 0.01]$. Thus, to minimize the number of *FP* without any significant increase in *FN*, the mismatch score threshold should be chosen in the range $[0.01, 0.02]$. As we increase n beyond 0.02, the number of *FN* would start increasing significantly without any substantial reduction in the number of *FP* due to which the detection performance would degrade. Thus, choosing n in the range $[0.01, 0.02]$ would result in the best detection performance of the scheme, as also apparent from Table 7.

7.2.2 Detection performance v/s different values of n


To analyze the effect of varying the window size n on the detection performance of the scheme, we experimented with four more time intervals of 8 hours each such that

TABLE 7: Detection performance

Interval	# of Flows		Mismatch score threshold	Window size n	Accuracy	FPR	Recall	Precision
	Benign	Attack						
1 st	12423	3242	0.01	5	98.55	1.83	100.00	93.46
2 nd	12423	3225	0.02	5	98.65	1.71	100.00	93.83
3 rd	12120	3222	0.03	5	93.39	1.64	074.71	92.36
4 th	10605	2860	0.04	5	88.91	1.64	053.88	89.85
5 th	10620	3070	0.05	5	82.80	1.62	006.08	43.23

TABLE 8: Effect of varying n on detection performance

Window size (n)	Accuracy	FPR	Recall	Precision
3	98.55	1.79	100	92.94
5	98.35	2.04	100	92.03
7	98.68	1.63	100	93.52
9	98.89	1.37	100	94.52

Fig. 6: Time required to extract lookahead pairs for different n

the mismatch score threshold was fixed, but n was varied across the intervals. We show the results of this experiment in Table 8. We can notice from the table that varying the window size does not virtually affect the proposed scheme's detection accuracy. However, we also observed from our experiments that the longer the chosen window size, the longer time was required to extract the lookahead pairs from an event sequence, as shown in Figure 6. This results in the delayed classification of an event sequence as a normal or anomalous sequence. Thus, choosing a smaller window size is essential so that the mismatch scores can be calculated and the corresponding decision can be made as soon as possible.

7.2.3 Detection latency and computational overhead

It is essential for an active detection scheme to detect anomalies as soon as possible so that an alarm can be raised and necessary preventive measures can be taken. Thus, we also measured how fast our proposed detection scheme could detect anomalous event sequences. Figure 7 shows the result of this measurement. We can notice from the figure that the




Fig. 7: CDF plot of the time required to detect anomalous event sequences

detection scheme could detect all anomalous sequences in virtually less than 30 seconds.

While the detection scheme was running at the Interpreter, we also monitored for 1.5 hours the computational overhead incurred due to its execution. In particular, we logged the computer's overall CPU⁵ utilization at regular intervals of 1 second using `psutil` and plotted the graph shown in Figure 8. We can notice from the figure that the CPU usage, while the detection scheme was running, was less than $\approx 7\%$ (dense blue area) in most of the time intervals, except for some intervals wherein the CPU usage reached up to $\approx 16\%$. Thus, due to the small computational overhead of the proposed detection scheme, it can be an ideal solution for detecting Slow HTTP/2 DoS attacks in real-time.

8 CONCLUSION AND FUTURE WORK

Slow Rate DoS attacks are a matter of grave concern for server administrators because of two prominent reasons. First, these attacks require significantly less computational power and thus, can be launched even from mobile devices. Second, these attacks generate minimal traffic, due to which they are highly stealthy. The effect of these attacks on HTTP/2 servers on the Internet had not been analyzed earlier. Moreover, the known mechanisms to counter the attacks against HTTP/1.1 can not counter them against HTTP/2. Thus, in this work, we attempted to bridge these gaps by

5. The Interpreter computer had AMD Ryzen 7 5800H processor.




Fig. 8: Computational overhead

first performing an empirical evaluation of Slow HTTP/2 DoS attacks on the Internet and subsequently proposing a scheme to detect these attacks in real-time. Our experiments showed that several HTTP/2 servers on the Internet **delay** closing the connections from which incomplete requests are sent, thus being vulnerable to the Slow Rate DoS attacks. Our experiments to test the performance of the proposed detection scheme showed that it could accurately detect the attacks in real-time with a marginal computational overhead. In the future, we plan to work on the vulnerability assessment of HTTP/3, the upcoming version of HTTP. As of January 2022, HTTP/3 has not been standardized, and it currently holds the status of an Internet Draft.

REFERENCES

- [1] N. Tripathi, N. Hubballi, and Y. Singh, "How Secure are Web Servers? An Empirical Study of Slow HTTP DoS Attacks and Detection," in *Proceedings of the International Conference on Availability, Reliability and Security (ARES)*, 2016, pp. 454–463.
- [2] N. Tripathi and N. Hubballi, "A Probabilistic Anomaly Detection Scheme to Detect DHCP Starvation Attacks," in *Proceedings of the International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, 2016, pp. 1–6.
- [3] —, "Preventing Time Synchronization in NTP Broadcast Mode," *Computers & Security*, vol. 102, p. 102135, 2021.
- [4] G. Mantas, N. Stakhanova, H. Gonzalez, H. H. Jazi, and A. A. Ghorbani, "Application-Layer Denial of Service Attacks: Taxonomy and Survey," *International Journal of Information and Computer Security*, vol. 7, no. 2/3/4, p. 216–239, 2015.
- [5] Gonzalez, Hugo and Gosselin-Lavigne, Marc Antoine and Stakhanova, Natalia and Ghorbani, Ali A, "The Impact of Application-Layer Denial-of-Service Attacks," *Case Studies in Secure Computing: Achievements and Trends*, vol. 261, 2014.
- [6] SANS ISC InfoSec Forums. (2009) Slowloris and Iranian DDoS Attacks. <https://isc.sans.edu/forums/diary/Slowloris+and+Iranian+DDoS+6622/>.
- [7] N. Tripathi and N. Hubballi, "Application Layer Denial-of-Service Attacks and Defense Mechanisms: A Survey," *ACM Computing Surveys (CSUR)*, vol. 54, no. 4, pp. 1–33, 2021.
- [8] M. Belshe, R. Peon, and M. Thomson, "(RFC 7540) Hypertext Transfer Protocol Version 2 (HTTP/2)," 2015.
- [9] N. Tripathi and N. Hubballi, "Slow Rate Denial of Service Attacks against HTTP/2 and Detection," *Computers & security*, vol. 72, pp. 255–272, 2018.

- [10] N. Tripathi and A. K. Shaji, "Defer No Time, Delays have Dangerous Ends: Slow HTTP/2 DoS Attacks into the Wild," in *Proceedings of the 14th IEEE International Conference on Communication Systems and NETWORKS (COMSNETS)*, 2022, pp. 194–198.
- [11] D. Beckett and S. Sezer, "HTTP/2 Cannon: Experimental Analysis on HTTP/1 and HTTP/2 Request Flood DDoS Attacks," in *Proceedings of the International Conference on Emerging Security Technologies (EST'17)*, 2017, pp. 108–113.
- [12] —, "HTTP/2 Tsunami: Investigating HTTP/2 Proxy Amplification DDoS attacks," in *Proceedings of the International Conference on Emerging Security Technologies (EST'17)*, 2017, pp. 128–133.
- [13] S. T. Zargar, J. Joshi, and D. Tipper, "A Survey of Defense Mechanisms against Distributed Denial of Service (DDoS) Flooding Attacks," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 4, pp. 2046–2069, 2013.
- [14] A. Praseed and P. S. Thilagam, "DDoS Attacks at the Application Layer: Challenges and Research Perspectives for Safeguarding Web Applications," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 661–685, 2019.
- [15] Core. (2019) <https://httpd.apache.org/docs/2.4/mod/core.html>.
- [16] mod_antiloris. (2013) <https://sourceforge.net/projects/mod-antiloris/>.
- [17] mod_limitipconn. (2002) <http://dominia.org/djao/limitipconn.html>.
- [18] mod_reqtimeout. (2019) https://httpd.apache.org/docs/trunk/mod/mod_reqtimeout.html.
- [19] Imperva (2016): HTTP/2: in-depth analysis of the top four flaws of the next generation web protocol. https://www.imperva.com/docs/Imperva_HII_HTTP2.pdf.
- [20] E. Adi, Z. A. Baig, P. Hingston, and C.-P. Lam, "Distributed Denial-of-Service Attacks against HTTP/2 Services," *Cluster Computing*, vol. 19, no. 1, pp. 79–86, 2016.
- [21] A. Praseed and P. S. Thilagam, "Modelling Behavioural Dynamics for Asymmetric Application Layer DDoS Detection," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 617–626, 2021.
- [22] —, "Multiplexed Asymmetric Attacks: Next-generation DDoS on HTTP/2 Servers," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1790–1800, 2019.
- [23] E. Cambiaso, M. Aiello, M. Mongelli, and I. Vaccari, "Detection and Classification of Slow DoS Attacks Targeting Network Servers," in *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES)*, 2020.
- [24] A. Praseed and P. S. Thilagam, "Fuzzy Request Set Modelling for Detecting Multiplexed Asymmetric DDoS Attacks on HTTP/2 servers," *Expert Systems with Applications*, vol. 186, p. 115697, 2021.
- [25] NETRESEC, "PolarProxy." [Online]. Available: <https://www.netresec.com/?page=PolarProxy>

ACKNOWLEDGMENTS

The authors would like to thank the Netresec team for supporting this research with a free L1 license of PolarProxy.



Nikhil Tripathi is currently working as an Assistant Professor in Computer Science & Engineering Group at the Indian Institute of Information Technology, Sri City, India. Prior to his current role, he worked as a cyber security researcher at Fraunhofer Institute for Secure Information Technology (SIT), Germany. He earned his Ph.D. in Computer Science and Engineering from the Indian Institute of Technology Indore, India, in 2019. He primarily works on the vulnerability assessment of popular application layer protocols to explore possible DoS vulnerabilities and devise novel defense mechanisms to counter large-scale DoS attacks. He has published research papers in various reputed journals such as *ACM Computing Surveys*, *Computers & Security*, and conferences like *INFOCOM*, *ARES*, and *COMSNETS*. He also holds professional memberships of societies such as *IEEE*, *IEEE Communications Society*, and *ACM*.