WAFFLED: Exploiting Parsing Discrepancies to Bypass Web Application Firewalls

Seyed Ali Akhavani *
Northeastern University

Bahruz Jabiyev
Northeastern University

Ben Kallus
Dartmouth College

Cem Topcuoglu
Northeastern University

Sergey Bratus
Dartmouth College

Engin Kirda
Northeastern University

arXiv:2503.10846v2 [cs.CR] 29 Apr 2025

Abstract

Web Application Firewalls (WAFs) have been introduced as essential and popular security gates that inspect incoming HTTP traffic to filter out malicious requests and provide defenses against a diverse array of web-based threats. Evading WAFs can compromise these defenses, potentially harming Internet users. In recent years, parsing discrepancies have plagued many entities in the communication path; however, their potential impact on WAF evasion and request smuggling remains largely unexplored. In this work, we present an innovative approach to bypassing WAFs by uncovering and exploiting parsing discrepancies through advanced fuzzing techniques. By targeting non-malicious components such as headers and segments of the body and using widely used content-types such as `application/json`, `multipart/form-data`, and `application/xml`, we identified and confirmed 1207 bypasses across 5 well-known WAFs, AWS, Azure, Cloud Armor, Cloudflare, and ModSecurity. To validate our findings, we conducted a study in the wild, revealing that more than 90% of websites accepted both `application/x-www-form-urlencoded` and `multipart/form-data` interchangeably, highlighting a significant vulnerability and the broad applicability of our bypass techniques. We have reported these vulnerabilities to the affected parties and received acknowledgments from all, as well as bug bounty rewards from some vendors. Further, to mitigate these vulnerabilities, we introduce `HTTP-Normalizer`, a robust proxy tool designed to rigorously validate HTTP requests against current RFC standards. Our results demonstrate its effectiveness in normalizing or blocking all bypass attempts presented in this work.

Keywords

Web Security, Web Application Firewalls, Fuzzing, HTTP, Parsing Discrepancies, Request Smuggling, WAF Evasion

1 Introduction

The widespread adoption of web applications has made them prime targets for cyberattacks. To protect these applications, Web Application Firewalls (WAFs) have been introduced as essential and popular security gates. These systems inspect incoming HTTP traffic to filter out malicious requests, and provide defenses against a diverse array of web-based threats, ranging from SQL injection to Cross-Site Scripting attacks, and beyond.

Despite their critical role, WAFs are not immune to evasion. Traditional WAF evasion techniques often rely on distorting attack payloads to bypass detection rules while ensuring the payloads remain executable by web applications. Attackers usually either obfuscate the payload with encoding schemes or inject new characters into payloads to bypass WAF rules. However, WAF vendors have already taken measures against most of these bypass techniques, which have been known for many years now. Also, these types of attacks assume the ability of the target web application to parse the encoded, or obfuscated payload.

As the threat landscape evolves, HTTP Request Smuggling (HRS) attacks have gained significant attention. HRS exploits discrepancies in the interpretation of HTTP requests between different entities in the communication chain, such as servers, proxies, and WAFs. These attacks can have severe consequences, including unauthorized access to sensitive information, session hijacking, and server compromise. The increasing complexity of web applications and their reliance on intermediary components have increased the risk of HRS vulnerabilities. Recent studies have investigated the importance of HRS, identifying new variants of attack and suggesting defense mechanisms [8, 13–15].

Parsing discrepancies, caused by inconsistencies in the interpretation of HTTP requests, play a critical role in enabling attacks such as HRS. These discrepancies first appeared in the communication path between servers, but WAFs, which are integrated to this path, may themselves be vulnerable to such inconsistencies. As a layer between the client and the web application, WAFs must correctly interpret HTTP requests to protect against malicious activities. However, vulnerabilities in their parsing mechanisms can allow attackers to exploit these discrepancies, bypassing the WAF, and allowing attacks to reach the web application. Building on this understanding of the limitations of current WAF defenses and the emerging threat of HRS, we present a novel, real-world approach to bypassing WAFs. Our method exploits content parsing discrepancies between WAFs and web application frameworks. Unlike traditional evasion tactics, we keep the attack payload intact, and focus on mutating specific content elements, such as the boundary in `multipart/form-data`, or namespace feature in `application/xml`, causing the WAF to misinterpret the content. This misinterpretation allows the payload to pass through, while the web application framework correctly parses and executes the attack.

Our work tested a wide range of combinations of popular WAFs, including Google Cloud Armor, Cloudflare, AWS WAF, Azure WAF, and ModSecurity on NGINX, alongside widely-used web application frameworks such as Flask, Laravel, FastAPI, Gin, Express, and Spring Boot. We focused on three complex content types: `multipart/form-data`, `application/xml`, and `application/json`. By repurposing a grammar-based and

* Corresponding Author.

structure-aware HTTP fuzzer, we identified implementation differences in how these WAFs and frameworks parse these content types. Our findings reveal that most WAF-framework pairs can be bypassed using various content distortions, highlighting a significant vulnerability in current WAF implementations.

In this work, we summarize our contributions as follows:

- We introduce a fundamentally new approach that uses content parsing discrepancies to bypass web application firewalls.
- We present a practical methodology for automatically finding new discrepancy-based bypass vectors using black-box fuzzing techniques.
- We design and implement a pipeline that tests these bypass vectors against popular web application firewalls and frameworks.
- We demonstrate successful discrepancy-based bypass instances on the pairs of most popular web application firewalls and frameworks, and we coordinate mitigation efforts with the affected technology vendors.
- We analyze the interchangeability of our bypass techniques using real-world data from PublicWWW, demonstrating that our findings are widely-applicable and practical across real-world web applications.
- We present the first tool, `HTTP-Normalizer`, that ensures that all proposed bypasses are avoidable by enforcing proper techniques.

Availability. WAFFLED, fuzzing grammars, WAF configuration, web application frameworks and parsers code, and `HTTP-Normalizer` are open source, and are available on this [1] repository.

2 Background

2.1 Web Content Types and RFCs Explained

Our work focuses on using specific features of content-types to bypass a malicious request; thus, it is crucial to inspect the details of the content and media types discussed in this paper.

RFCs (Request for Comments). RFCs [10] are a series of documents that define protocols, procedures, and conventions used on the Internet and networking standards. They are published by the Internet Engineering Task Force (IETF) and related organizations. Each RFC is assigned a unique number, and these documents serve as the authoritative source of information on various Internet standards and protocols. For instance, the structure of HTTP headers, MIME types, and various other web technologies are defined by respective RFCs.

ABNF (Augmented Backus-Naur Form). ABNF grammar rules [5] are a formal notation used to specify the syntax of RFCs. ABNF extends the basic BNF notation to provide a more flexible and precise way to define the syntax of Internet protocols. It is particularly important in our work as it allows for precise definitions of content types and headers, which we exploit to identify parsing discrepancies in WAFs. By understanding and manipulating ABNF rules, we can generate requests that challenge the ability of WAFs to consistently parse and enforce these standards.

Multipart. Multipart content-types allow HTTP messages to include multiple entities with varying media types in a single message body. They are crucial for handling file uploads and complex data structure support. These content-types, such as multipart/form-data, multipart/related, and multipart/mixed, are governed by specific RFCs (e.g., RFC 2387, RFC 7578, RFC 2045-2047) [6, 7, 16–18] that define their structure and usage scenarios. Our research investigates the complexities of multipart/form-data requests, exploiting features such as boundary definitions, charset, content-disposition, and other custom headers.

XML. Extensible Markup Language serves as a flexible format for structured data representation, commonly used in data exchange protocols. XML-based requests, defined in RFC 7303 [21], are characterized by elements such as DOCTYPE declarations, schemas, and CDATA sections. Our approach includes mutating these structural elements to investigate WAF responses to XML-specific parsing matters. For example, variations in DOCTYPE declarations and CDATA usage are tested to identify vulnerabilities in WAF handling of XML content.

JSON. JavaScript Object Notation, described in RFC 8259 [4], is a lightweight data interchange format widely used in modern web applications. JSON requests are characterized by their simple structure based on key-value pairs. We manipulate JSON object formatting and nested structures to assess WAF handling of JSON parsing anomalies.

2.2 How WAFs Work

To understand the functioning of WAFs, one must visualize the transactional exchange between web applications and users. Before a web application receives an HTTP request, the WAF intercepts this request, serving as a firewall positioned between the application and the user. It thoroughly scrutinizes the request by analyzing its parameters, headers, and payload to identify any potential signs of malicious intent. If the request appears legitimate, the WAF allows it to pass through to the web application, enabling normal operation. However, if the request exhibits characteristics of an attack, the WAF either completely blocks the whole request, or removes the malicious part of it, preventing the malicious code from executing and protecting the web application from harm. Figure 1 represents this process. There are many industrial and open source WAFs available to the public, each using their own techniques to prevent malicious requests to be passed to the protected web application. Notable WAF providers include Cloudflare, Google Cloud Armor, Microsoft Azure WAF, Amazon AWS WAF, and ModSecurity.

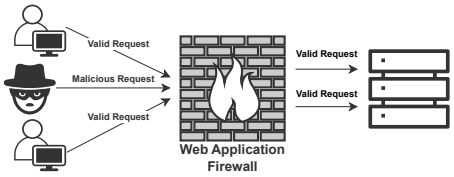

Figure 1: WAF filtering malicious requests.

[1] https://github.com/sa-akhavani/waffled/

2.3 Foundations of WAF Operations and Ruleset Definitions

A WAF operates by inspecting HTTP traffic between a client and a web application, filtering out harmful requests to prevent attacks such as SQL injection, cross-site scripting (XSS), and malicious script uploads. The WAF is positioned inline with the traffic flow, allowing it to analyze and intercept malicious content before it reaches the application server. WAFs employ several detection approaches.

Managed Rulesets. WAFs use predefined sets of rules that define patterns associated with common attacks. These rulesets are frequently updated to respond to new threats. Most WAFs use their own custom rulesets. But one of the famous rulesets that is widely used among almost all WAFs and most custom rulesets use it as a base, is OWASP CRS (Core Rule Set) [3]. This ruleset is a set of generic attack detection rules for use with compatible WAFs. It aims to protect web applications from a wide range of attacks, including the OWASP Top Ten. CRS provides protection against many common attack categories, including SQL Injection, Cross Site Scripting, Local File Inclusion, etc.

Signature-Based Detection. This method involves identifying known attack vectors based on specific signatures or patterns in the data. It is effective against known threats, but can be limited in identifying novel or modified attacks.

Machine Learning-Based Detections. Some advanced WAFs leverage machine learning algorithms to detect anomalies in traffic patterns. These models can adapt to new threats over time, offering a dynamic defense mechanism.

A WAF operates by scanning incoming HTTP requests, analyzing headers, cookies, URL parameters, and body content. Each request is parsed to extract these fields, allowing the WAF to evaluate them against a predefined set of rules. These rules follow a structured format:

```
if <field><operator><value>
then <action>
```

Listing 1: WAF Rule Format

- **Field**: This could be any part of the HTTP request such as a header, cookie, MIME type, or URL parameter.
- **Operator**: Common operators include *equals*, *contains*, *not equal*, etc.
- **Value**: The specific value to match against.
- **Action**: The action to be taken if the rule matches, such as *block*, *skip*, or *send security challenge*.

To accommodate complex attack patterns, WAF rules often rely on regular expressions (regex). For instance, detecting SQL injection attempts may involve scanning URL parameters or request bodies for SQL-related keywords. When a request arrives, the WAF extracts relevant fields, evaluates them against the defined rules, and executes the appropriate action. If a match is found, the WAF may block the request, log the event, or challenge the user with additional security measures such as a CAPTCHA. In cases where no match occurs, the request is forwarded to the application without interference.

Listing 2 presents an example of how a WAF rule might be defined and applied. First, there should be a rule definition.

```
if url_parameter "user_input" contains "DROP TABLE"
then block
```

Listing 2: Example WAF Rule. Field: url_parameter *user_input*, Operator: contains. Value: DROP TABLE, and Action: block

After a request is received by a WAF, it does request Parsing and Rule Application. If a request is received with the URL parameter: `user_input=DROP TABLE users`, the WAF extracts the `user_input` parameter, and evaluates it against the rule. `if "DROP TABLE" in user_input:` The regex "DROP TABLE" matches the parameter value. Since the action for such pattern is defined as "block", the WAF executes the block action, denying the request and logging the incident.

By using well-defined rules and regular expressions, a WAF effectively protects web applications from a wide range of attacks. The flexibility in defining rules allows security teams to tailor the WAF to the specific needs of their applications, enhancing the security without significantly impacting performance.

As explained, a critical aspect of WAF operation is how it parses and interprets incoming data. Parsing behaviors can differ significantly across various WAFs, leading to discrepancies in how they process the same set of data. These discrepancies present a potential avenue for bypassing WAFs, as attackers can exploit differences in parsing to craft requests that are interpreted harmlessly by the WAF, but maliciously by the target application. This work primarily focuses on uncovering new bypassing techniques by abusing these parsing discrepancies.

3 Related Work

3.1 Architecture, Payload, and Protocol Level Bypasses.

At the architectural level, bypasses exploit vulnerabilities within the web application's infrastructure. Techniques such as IP spoofing and server-side request forgery (SSRF) allow attackers to access the origin server directly, circumventing WAF protection. Cache poisoning attacks, as discussed by Meiners et al. [19], exploit CDN vulnerabilities to disrupt access to web applications. Additionally, file processing exploits, such as those highlighted by Freiling et al. [12], offer insights into similar architectural vulnerabilities.

Bypassing WAFs at the protocol level typically involves manipulating communication protocols. HTTP request smuggling, where a malicious request is embedded within a legitimate one, is a key example, with T-Reqs [11] providing methods to identify such vulnerabilities. Similarly, features such as chunked transfer encoding, which involves fragmenting HTTP data into smaller chunks, can be exploited to evade WAF detection. Dalili's work on IP fragmentation [1] explore these protocol-level evasion strategies. Also, SYMTCP [24] focuses on evading deep packet inspection, a principle that can be applied to circumvent WAFs at the protocol level. Furthermore, Handley et al. [9] inspects the evasion of network

intrusion detection systems, a concept closely related to bypassing WAFs at the protocol level.

Payload level bypassing techniques involve transforming the original payload to render it undetectable by the WAF. Obfuscation techniques, such as encoding or encryption, are frequently employed to conceal malicious code from WAF scrutiny. Additionally, evasion techniques, such as altering the order of characters or incorporating whitespace, can be used to confuse the WAF, and facilitate bypassing. Payload size manipulation, such as oversized POST requests that exploited a vulnerability in Google Cloud Armor [2], illustrates how payload modifications can defeat WAF protections.

Not all bypasses are feasible using by targeting a single level approach above. Some works involve combining multiple techniques to achieve a bypass. AutoSpear [20] exemplifies a tool that automates WAF bypassing using a combination of architectural, protocol, and payload-level techniques. Our work also fits in this category, since we use both features of protocol-level and payload-level techniques, although not modifying the attack payload.

The academic works discussed here significantly enhance our knowledge of bypass techniques and WAF behavior, providing a platform for further exploration in WAF bypassing methods. Our work builds upon this foundation by introducing a novel approach that exploits content-parsing discrepancies, pushing the boundaries of existing methodologies. Unlike prior research, which often focuses on obfuscating payloads or manipulating protocol features, our approach targets fundamental weaknesses in how WAFs parse HTTP requests.

3.2 Differences with Prior Work on WAF Bypasses

In recent work, Wang et al. [23], presented WAFManis, in which they examined bypassing WAFs using protocol-level evasion. While this work aligns with our work in exploring WAF vulnerabilities, our research fundamentally differs in several key aspects.

Scope. While previous work has primarily concentrated on protocol-level attacks, our work uniquely targets content-type-specific attacks by exploiting parsing discrepancies within HTTP requests. That is, our approach extends beyond the protocol layer to focus on inconsistencies in both headers and bodies, targeting a wider array of bypass scenarios. Although there is some overlap in our findings with Wang et al., particularly concerning request header mutations, this is not the primary focus of our research. In contrast, by systematically examining all web-used content-types, our approach offers a broader scope, potentially identifying bypass techniques that have not been previously explored. Both works advocate for a systematic approach, yet our focus on content-type-specific parsing discrepancies enables the identification of a novel category of WAF bypasses.

Methodological Differences and Expanded Coverage. Although Wang et al. employ fuzzing to uncover WAF bypasses, the methodologies diverge significantly. Our approach emphasizes content-type-specific RFC rules and features within request bodies, such as `application/xml` namespaces, `multipart/form-data` custom headers and charsets, and `application/json` structures. This distinction introduces a systematic exploration of bypasses through various content-types, leading to the discovery of new categories of bypasses beyond what is covered in previous work.

Classification. We propose a classification system for bypass techniques, dividing them into 24 distinct types based on content types, which facilitates easier detection, analysis, and reporting. In contrast, Wang et al. defines only 3 attack types, which lack the granularity needed for in-depth study and mitigation. Our system also enables the identification of potential defenses for each bypass, as discussed in the evaluation.

Practicality. Our research demonstrates the real-world applicability of our findings by testing bypass techniques on 100 live websites, while Wang et al.'s work was limited to a controlled environment, highlighting the broader relevance of our work.

Mitigation. We introduce the `HTTP-Normalizer`, a proof-of-concept strategy designed to eliminate parsing discrepancies by enforcing strict RFC compliance, a contribution not present in previous work.

4 WAFFLED Design and Methodology

Figure 2 provides an overview of the WAFFLED approach for finding discrepancy-based bypass vectors for WAFs. The methodology involves generating, mutating, prefiltering, and testing HTTP requests across the WAFs and web application frameworks to identify successful bypass instances.

4.1 Input Generation and Mutation

The fuzzer generates valid HTTP requests using a predefined grammar that contains a web attack payload (e.g., SQL injection, Cross-Site Scripting). Mutations are applied everywhere except the attack payload itself, ensuring that bypass attempts focus on discrepancies rather than payload obfuscation. The highlighted part in Listing 3 shows an XSS attack payload, which remains unmodified during mutations.

These mutations are expected to give us new discrepancy-based bypass instances. The goal is to confuse WAFs so they fail to properly parse the mutated requests and allow them through, while the web application frameworks, despite the mutations, correctly interpret and execute the attack payload.

4.2 Identifying Successful Bypass Instances

The target infrastructure consists of all tested WAFs where each WAF is able to forward requests to all web application framework instances. Each instance is a tiny web application that parses received requests by using the parser library of the web application framework (e.g., Flask) for the tested content type (e.g., `multipart/form-data`). The results of parsing operation are logged for a later analysis. If a log shows that the request is parsed successfully and the request body contains the attack payload, then the corresponding input request can be used to successfully bypass the relevant WAF.

4.3 Motivating Example

Listing 4 presents an example of how parsing discrepancies can lead to WAF circumvention. This example has two boundaries defined: `fake-boundary` and `real-boundary`. When this request arrives at

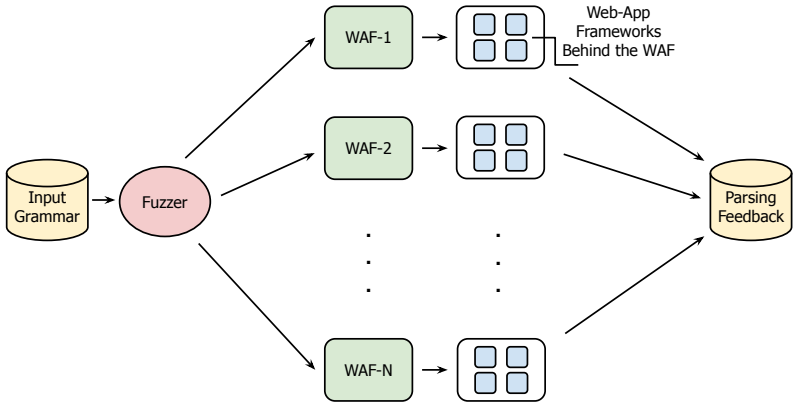

Figure 2: WAFFLED Overview.

the WAF, it picks the first one (i.e., `fake-boundary`) as the boundary value and ignores the remaining boundaries, thereby missing the attack payload. When it is forwarded to the web application framework running behind the WAF, as it supports the boundary continuation mechanism – which was standardized in RFC 2231 and allows the use of multiple parameters to contain a single parameter value – it takes the boundary value to be the concatenation of those two values (i.e., `real-boundary`), thereby parsing the XSS payload encapsulated between the `real-boundary` boundaries.

```
POST / HTTP/1.1
Host: victim.com
Content-Length: 114
Content-Type: multipart/form-data; boundary=1234

--1234
Content-Disposition: form-data; name="field1"

<script>alert(document.cookie)</script>
--1234--
```

Listing 3: A malicious HTTP request using multipart/form-data content type with an XSS attack payload. The malicious part of the payload is highlighted and unmodified during the mutations.

```
POST / HTTP/1.1
Host: victim.com
Content-Length: 230
Content-Type: multipart/form-data;
boundary=fake-boundary;boundary*0=real-;boundary*1=boundary

--fake-boundary
Content-Disposition: form-data; name="field1"

value1
--fake-boundary--
--real-boundary
Content-Disposition: form-data: name="id"

<script>alert(document.cookie)</script>
--real-boundary--
```

Listing 4: A malicious `multipart/form-data` request that uses parameter continuation.

5 Experiment Setup

5.1 Tools and Infrastructure for Experiments

We used a modified version of the T-Reqs [11] fuzzer for the experiments. The modifications include adding support for encrypted HTTP requests, and setting `content-length` dynamically based on the length of the body. We also modified the T-Reqs source code to avoid using < and > characters to parse the grammar, since these characters were also used in the XSS attack payload that was present in generated requests.

Input grammars were created in accordance with RFC specifications, encompassing all standard components of the content type and their possible values. The fuzzer was executed using these specific grammars[2] for several hours, generating and testing a total of 373,670 requests across 5 WAFs and 6 web application frameworks.

Our server infrastructure was deployed across multiple cloud platforms, including Amazon AWS, Google Cloud, and Microsoft Azure. This configuration was utilized to host our web applications and conduct the experiments.

5.2 WAF Configuration and Setup

To comprehensively assess the resilience of WAFs against mutated HTTP/1.1 web requests, we conducted an extensive evaluation of several WAFs. The selection criteria focused on solutions with significant adoption across both commercial and open-source domains. The WAFs included in this study were AWS WAF, Cloudflare WAF, Google Cloud Armor, Microsoft Azure WAF, and ModSecurity on NGINX. This selection enabled us to evaluate the effectiveness of both widely-used community-driven tools and leading commercial solutions in the industry.

For all tested WAFs, default settings were utilized without any modifications to preserve the integrity of the testing environment and ensure the results accurately reflected typical user experiences. The consistent configuration, especially with the widespread adoption of the OWASP CRS, enabled a fair comparison across WAFs despite their differing underlying technologies and implementation

[2]Available in the public repository

Table 1: List of Tested WAFs and Their Supported Request Content-Types (As of July 2024).

WAF	Content-Type	Bypassed
Cloudflare WAF	multipart, json, xml	Yes
Microsoft Azure WAF	multipart, json	Yes
Google Cloud Armor	multipart, json	Yes
Amazon AWS WAF	multipart, json	No
ModSecurity on NGINX	multipart, json, xml	Yes

strategies. This standardized methodology established a clear baseline to evaluate each WAF's capability to process mutated HTTP requests, effectively highlighting their individual strengths and potential vulnerabilities.

Table 1 lists the WAFs and their supported content-types in our study.

AWS WAF. AWS WAF was configured using Elastic Load Balancing (ELB). We employed the AWS Managed Ruleset in conjunction with OWASP CRS version 3.0, using default settings to accurately represent typical user configurations.

Cloudflare WAF. Cloudflare WAF was set up as the DNS provider for our domain, applying its WAF to all incoming traffic via its global CDN. We used the Pro version, enabling both Cloudflare's Managed Ruleset and the OWASP CRS to ensure comprehensive threat coverage.

Google Cloud Armor. Google Cloud Armor was deployed on a virtual machine behind Google Cloud's Load Balancer. We configured it with rulesets for SQL injection and XSS protection (`sqli-v33-stable` and `xss-v33-stable`) at Sensitivity Level 1.

Microsoft Azure WAF. Microsoft Azure WAF was set up through Azure's Application Gateway using the WAF V2 tier, with OWASP CRS version 3.0 enabled by default. This configuration aligns with common deployment practices.

ModSecurity. ModSecurity was deployed as a plugin for NGINX, configured with OWASP CRS version 3.0 at Paranoia Level 1, reflecting the default settings commonly used in many environments. The setup adhered to the official documentation guidelines, ensuring an accurate evaluation of its effectiveness in scenarios typical of real-world deployments.

The influence of rulesets on our findings is minimal for several key reasons. First, the attack payloads used in this study are intentionally straightforward, designed to be blocked by all standard rulesets across the tested WAFs, as confirmed during preliminary testing. Second, our bypass techniques do not depend on obfuscating the payload itself, but rather exploit parsing inconsistencies, focusing on how WAFs interpret HTTP requests before applying their rulesets. This approach highlights that the bypasses stem primarily from the WAF's inability to correctly parse the request or its decision to bypass parsing due to the complexity of the input, irrespective of the specific ruleset in use.

5.3 Web Application Frameworks Setup

In this work, we evaluated seven widely-used web application frameworks: Express, Node.js-HTTP, Flask, FastAPI, Gin, Laravel, and Spring Boot. To comprehensively analyze their parsing capabilities, we employed a combination of the frameworks' default parsers and HTTP parser packages sourced from reputable platforms. The selection of parser packages was informed by their prevalence and popularity within the developer community, as indicated by metrics from resources such as npm and PyPi Stats. This approach ensured that our study included modules that are both widely adopted and representative of real-world usage.

To accurately test how each framework and parser handles the request body, we implemented the following approach: we attempted to access the parsed content of the request body through framework-specific methods such as `request.body.parse()`, etc. The goal was to retrieve the value of a specific field (`field`), and have the framework and parser process the entire request body. After parsing, we searched for the sent attack payload within the parsed content. If any of the parsed fields contained the attack payload, it indicated that the request successfully bypassed the WAF, and was correctly parsed by the framework and parser. We then sent a success flag along with instance name and WAF information back to our client to mark the successful attacks.

For parsing request bodies in our study, each framework's default parser was utilized. In cases where the default parser did not support a specific content-type, we employed the most popular third-party parsers for that content-type. This approach of using both default and third-party parsers provided an understanding of how web application frameworks handle and interpret HTTP requests under different parsing scenarios.

For Node.js, the chosen parsers included `Busboy` and `Formidable` for parsing multipart content types, and `fast-xml-parser` for handling `application/xml` requests. Since Express and the HTTP module do not support parsing of multipart and xml content-types. For Python frameworks, we integrated `xmltodict` and `xmlminidom` for XML parsing. Multipart content types in Python were processed using `python-multipart` parsers.

In all our tests, we strictly adhered to the official documentation and examples of each framework, avoiding any custom development. This ensured that our evaluation covered the most common and correct usage patterns of these frameworks and parsers.

5.4 Validation of Bypass Effectiveness

To validate WAF bypasses with malicious payloads, we employed a rigorous methodology. Each test began with the dispatch of an initial request, embedded with a malicious payload, through the WAF. Upon receipt by the web application framework, the framework attempted to parse the request and extract the content of the body, subsequently storing it in a designated file on the server.

Our validation process compared the original malicious payload with the content received and stored on the target server. A successful bypass was determined if these two elements – the initial payload and the stored content – were found to be identical. We developed a script that compared all input payloads with their corresponding stored versions on the server. This approach allowed us to precisely determine whether the malicious payload had circumvented the WAF and was accurately parsed by the target web framework.

Table 2: List of tested web frameworks and whether they are vulnerable to at least one bypass in our study or not.

Framework	Version	Language	Vulnerable
Laravel	10.48.16	PHP	Yes
Spring Boot	3.2.2	Java	Yes
Gin	v1.9.1	Go	Yes
Express	4.18.2	Node.js	Yes
Fastapi	0.109.2	Python	Yes
Flask	3.0.2	Python	Yes
Node.js-HTTP	18.16.1	Node.js	Yes

This process was necessary because defining a successful bypass solely based on the ability of a request to evade the WAF would lead to misclassification. Such a definition could include malformed requests-those not conforming to HTTP standards or unprocessable by any framework-as successful bypasses, inflating false positives and obscuring real-world risks due to:

Unrealistic Attack Scenarios. Most WAFs allow malformed requests to pass because they cannot parse them. If we were to consider these instances as bypasses, more than 50% of mutated requests could be misclassified as successful.

Practical Mitigation Challenges. Malformed requests are frequently rejected at the framework level. Reporting these would inflating false positives and obscuring actionable findings.

Real-World Impact. A true attack occurs when a malicious request bypasses the WAF, is parsed by the framework, and the payload is executed. Malformed requests that frameworks cannot parse pose lower security risks and are lower priority.

All bypasses in our work represent successful attacks, meaning the request bypassed the WAF and was parsed successfully by a web application framework, enabling the malicious payload to be executed (e.g., XSS or SQL injection).

5.5 Experiment Results and Analysis

Table 2 provides the list of frameworks and Table 7 in the Appendix A provides the list of parsers that were inspected in our work. The experiments revealed that all of the frameworks and parsers were vulnerable to at least one bypass technique, highlighting the effectiveness and practical implications of our findings.

6 Findings

Before diving into analyzing these bypasses and reporting them, we first, analyze all results to retain only those bypassed requests that are achieved with the minimum number of mutations. Then, we classify bypasses based on their mutations. This allows us to:

- Identify the minimum mutations needed for a request to bypass a WAF uniquely.
- Ensure our results are realistic, and our reports are minimal and accurate by focusing solely on unique bypasses.
- Reporting and documenting bypass techniques and ensuring that mitigation strategies can address entire classes of bypasses rather than individual instances.

6.1 Identification of Unique Bypass Requests

In the process of testing and validating bypass techniques against WAFs, identifying unique bypass requests is crucial for accurately assessing the effectiveness of the discovered vulnerabilities. During our experiments, we discovered 1207 mutated requests that were successful bypasses. The identification of unique bypass requests involves rigorous analysis and filtering of the mutated HTTP requests generated during the testing phase.

Not all successfully bypassed requests are unique in terms of the underlying techniques used. For instance, consider a scenario where two mutations contribute to the bypass success:

- **Mutation I**: Introduces a `\x00` character into the boundary value of a multipart request, transforming `-boundary` to `-\x00boundary`.
- **Mutation II**: Changes the capitalization of characters in the `charset` parameter of the `Content-Type` header from `utf-8` to `Utf-8`.

While Mutation I, alone, constitutes a successful bypass strategy, Mutation II does not contribute to bypassing the WAF independently. However, when both mutations are combined into a single request, the request successfully evades WAF detection due to the presence of Mutation A. In this scenario, the unique bypass is attributed to Mutation A, whereas Mutation B does not qualify as an independent bypass strategy. And the request that is formed by both of these mutations is not a unique bypass.

To maintain the integrity and accuracy of the findings, redundant bypass requests that do not introduce unique evasion techniques are filtered out during the analysis phase.

After applying this minimization technique, among all of the generated requests, 1207 of them were unique bypasses found for (WAF, Framework) pairs.

6.2 Classification of Bypasses

After extracting the bypasses with the minimum number of mutations, we proceed with the classification of these bypasses. This classification process is essential as it allows us to systematically analyze and categorize the bypasses, facilitating the development of effective defense mechanisms.

To achieve this, each bypass is examined based on the unique mutation strategy and the mutated element used rather than the specific mutation. For instance, consider a JSON request with the body `{"field1": "value1"}`. If two distinct bypasses are achieved by inserting a `\x00` or `\x02` after the second double quote, both are classified under the same category: *manipulating the field name wrapper*. Although different characters are used, the fundamental concept remains consistent. This categorization simplifies the study of bypass techniques by grouping similar strategies, thereby enabling a more structured approach to understanding and mitigating these vulnerabilities.

6.3 Bypassed Requests Analysis

We now present the classification results of the bypassed requests for each content-type in our work. Some of the mutation classes are common among all tested content-types, they are defined once but mentioned for each content-type that they were leading to a successful bypass.

Figure 3: Occurrence of Successful Bypasses For Multipart Content-Types in Tested Frameworks

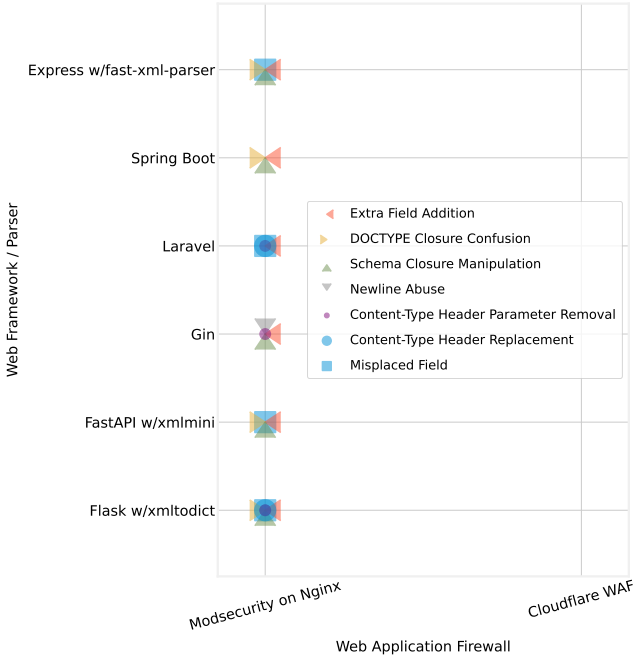

Figure 4: Occurrence of Successful Bypasses For application/xml Content-Types in Tested Frameworks

6.3.1 Multipart Bypass Classes. Our examination identified a total of 351 unique bypasses related to multipart content-type parsing. Figure 3 demonstrates valid bypass classes that we found in our work for multipart content-type. Table 3 contains examples for each discussed bypass category.

Boundary Delimiter Manipulation. This category involves the removal or alteration of boundary delimiters within multipart content. For instance, removing the `\r\n` sequence before the boundary string can confuse the parsing logic, leading to successful bypasses.

Content-Type Parameter Tweak. This technique includes modifications to the global Content-Type header's name, such as removing or inserting characters.

Content-Disposition Disruption. This technique targets the Content-Disposition header within the multipart content, altering its structure to evade detection.

Disrupted Header Injection to Body. This category includes adding redundant headers with disrupted header name directly into the multipart body content, which can mislead the WAF into processing the content incorrectly.

Content-Type Tweak in Body. Here, manipulations are performed on the Content-Type value within the body and not in the global header, such as inserting characters.

Charset Value Alteration in Body. This involves altering the charset value within the body content, affecting how the WAF interprets the encoding of the payload.

Header Separator Manipulation in Body. This technique modifies the separator between multiple header lines in the multipart body content, such as replacing newlines with other characters.

Boundary Delimiter Removal. This category involves the complete removal of boundary delimiters, affecting the WAF's ability to correctly parse multipart sections.

Linefeed Removal. This technique removes newline and carriage return characters. For example, removing the linefeed after the request headers and before beginning the boundary delimiter for the multipart request body. This approach can cause the WAF to misinterpret the structure of the multipart content.

Boundary Header Tampering. In this category, manipulations involve changes to the boundary value in the header, such as appending a semicolon to the boundary parameter in the content-type header line.

Whitespace Alteration. In this category, the whitespace character is replaced with `\t` in the content-type header. The request header must use the parameter value continuation feature of HTTP to define the boundary for this bypass to be effective. `boundary*0=re;boundary*1=al`

Disrupted Body Field. In this category, manipulations involve inserting invalid characters in a field name. The request header must use the parameter value continuation feature of HTTP to define the boundary for this bypass to be effective.

6.3.2 XML Bypass Classes. Our analysis revealed a total of 299 unique bypasses for XML content types. Figure 4 demonstrates valid bypass classes for application/xml content-type. Table 4 contains examples for each discussed bypass category.

Extra Field Addition. This category involves adding an extra field outside the defined XML schema. For example, inserting a new field such as `<field2>value2</field2>` exploits schema validation weaknesses, allowing malicious payloads to bypass security checks.

Table 3: Classification of Multipart Bypass Categories with Examples. Removals are displayed with a strike-through text with a gray background, while additions and replacements are shown with only a gray background.

Category Name	Request Example
Boundary Delimiter Manipulation	~~\r\n~~ –boundary
Content-Disposition Disruption	content-disposition: form-da \x00 a;
Distorted Header Injection to Body	conten\x00-extra: something
Content-Type Tweak in Body	Content-Type: text/plain \x00 ; charset=UTF-8
Charset Value Alteration in Body	charset= \x00 UTF-8
Header Separator Manipulation in Body	content-disposition: form-data; name="f1" \x00
Content-Type Parameter Tweak	C ~~o~~ ntent-Type: multipart/form-data;
Boundary Delimiter Removal	~~–boundary~~
Linefeed Removal	Content-Type: multipart/form-data; boundary=real\r ~~\n~~ \r\n
Whitespace Alteration	Content-Type: \t multipart/form-data; boundary*0=re;boundary*1=al
Disrupted Body Field	content-disposition: form-data; name="field1 \x00 "
Boundary Header Tampering	–boundary=value ;

DOCTYPE Closure Confusion. This class confuses the WAF by placing an extra character at the end of the xml body which confuses the WAF in parsing the DOCTYPE entity of the XML.

Schema Closure Manipulation. This technique manipulates the closure of XML schemas. By inserting characters, new elements, or duplicated field names and values at specific positions within the schema, the structure of the XML document is altered in a way that evades detection by the WAF.

Newline Abuse. This method achieves the bypass by placing an extra new-line before the content-type header.

Content-Type Header Parameter Removal. Bypasses achieved by removing the Content-Type header's name. Such manipulations exploit the dependency of some WAFs on specific header configurations to enforce security rules.

Content-Type Header Replacement. In this technique, the parameter name within the Content-Type header is replaced with its value. This manipulation can confuse WAFs that rely on precise header structures for detection.

Misplaced Field. This method involves placing field values outside their corresponding XML tags. By altering the position of values within the XML document, the attack payload can bypass WAF rules that expect a specific structure.

6.3.3 JSON Bypass Classes. In our work, we identified a total of 557 unique bypasses for JSON content types. Figure 5 demonstrates valid bypass classes for application/json content-type. Table 5 contains examples for each discussed bypass category.

Content-Type Removal. This category involves bypassing the WAF by removing the `Content-Type: application/json` header from the global header. This manipulation exploits the reliance of some WAFs on specific content-type headers to enforce security rules.

Field Wrapper Manipulation. In this category, the attack is achieved by adding characters such as `\x00` between the field name and the colon in the JSON object. This manipulation disrupts the field parsing.

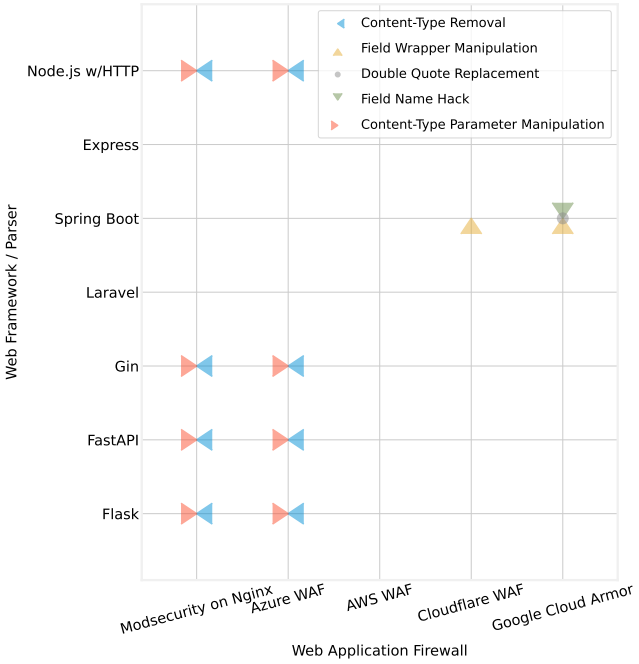

Figure 5: Occurrence of Successful Bypasses For application/json Content-Types in Tested Frameworks

Double Quote Replacement. This method replaces the double quotes surrounding the field names and values with other characters.

Field Name Hack. This technique involves altering characters within the field names of the JSON object. For example, replacing a character in the field name with `"\x00"` can bypass WAFs that do not properly handle such anomalies in field names.

Content-Type Parameter Manipulation. Here, the manipulation targets the Content-Type header name.

Table 4: Classification of XML Bypass Categories with Examples

Category Name	Request Example
Extra Field Addition	`<field1>value1</field1> <field2 attr="history">hi</field2>`
DOCTYPE Closure Confusion	`<!DOCTYPE BOOK [...]><field1>value1</field1> ] </BOOK>`
Schema Closure Manipulation	`<genre:schema><field1>value1</field1> j </genre:schema>`
Newline Abuse	`\r\n Content-Type: application/xml`
Content-Type Header Parameter Removal	~~`Content-Type:`~~ `application/xml`
Content-Type Header Replacement	`application/xml application/xml`
Misplaced Field	`... value1 <field1>value1</field1>...`

Table 5: Classification of JSON Bypass Categories with Examples.

Category Name	Bypassed Request Example
Content-Type Removal	~~`Content-Type: application/json`~~
Field Wrapper Manipulation	`{ "field1" \x00 : "<script>alert(document.cookie)</script>" }`
Double Quote Replacement	`{ "field1 \x00 : "<script>alert(document.cookie)</script>" }`
Field Name Hack	`{ "f \x00 eld1": "<script>alert(document.cookie)</script>" }`
Content-Type Parameter Manipulation	`Content - Type: application/json`

6.4 Practicality of Findings

The bypass findings of this paper for three content types are practical and useful only if these content types are popular amongst the web applications in the real world. We performed a study on a sample collected from real-world applications to gauge the popularity of these content types. In addition, we performed another study to check interchangeability between the content-types which can also contribute to the practicality of our findings.

We collected 100 high-ranking websites from PublicWWW [3], which is a search engine for website source codes, by using the query below which was tailored for the needs of the study.

```
type="email" "forgot password" -reCAPTCHA
```

This search query allowed us to find "forgot-password" web pages which contain an HTML form with an email-type input field and does not contain a CAPTCHA. The presence of CAPTCHA does not allow replaying the "forgot-password" requests, which makes it easier and faster to check the interchangeability of content types. We chose the first 100 search results, which are the "forgot-password" pages of popular websites ranking between 3K-50K. About half of these webpages were excluded from the study for reasons such as "non-english language", "obscenity" and "an emerging captcha".

On each "forgot-password" page, we submitted the form after typing in a non-existent email address and intercepted the generated HTTP request on the Burp Suite tool. We examined the `Content-Type` header field and the request body format to decide the content-type for each request. We found that more than two-thirds of the websites use the `application/x-www-form-urlencoded`, while one-fourth of them use the `application/json` and its variation. Only two websites use `multipart/form-data` and one uses URL parameters.

For each intercepted "forgot-password" request, to check the interchangeability between content types, we used the "Change body encoding" feature of Burp Suite (in the "Repeater" section) to convert the content type to `multipart/form-data` in `application/x-www-form-urlencoded` requests. We then examined the response status code and response body to check the sameness of responses, which indicates that the application does not differentiate those two content types. We found that more than ninety percent of websites accept them both.

These results suggest that our bypass findings are widely applicable and practical across web applications. The largest portion (i.e., more than two-thirds), which uses `application/x-www-form-urlencoded`, are affected by our `multipart/form-data` findings, since an attacker can easily switch the content type and apply the bypass techniques. The second largest portion, which uses `application/json`, is affected by all the relevant bypass findings we reveal in this paper. The fact that these two content types together make up more than 90% of content types used across web applications, clearly shows the extent of the practicality and impact of these bypass techniques.

7 Protecting WAFs and Frameworks by Normalizing HTTP Request Bodies

In this section, we discuss our approach to mitigate the identified vulnerabilities by introducing an HTTP multipart body normalizer that protects web application firewalls (WAFs) and web application frameworks from parsing-related attacks. The normalizer acts as a gateway for the WAF that strictly enforces RFC grammar rules

[3] https://publicwww.com/

```
POST / HTTP/1.1
Host: target.com
Content-Type: multipart/FoRm-dAtA; boundary="1234"
Content-Length: 90

--1234
Content-DISPOSITION:\tform-data;name="files";\t filename
     ="ab.txt"

Foo
--1234--
```

Listing 5: A request before normalization.

```
POST / HTTP/1.1
Host: target.com
Content-Type: multipart/form-data; boundary=1234
Content-Length: 90
Accept: */*
Accept-Encoding: gzip, deflate
User-Agent: Python/3.12 aiohttp/3.9.5

--1234
Content-Type: text/plain
Content-Disposition: form-data; name="files"; filename="
     ab.txt"

Foo
--1234--
```

Listing 6: The same request after normalization. Note the standardized capitalization and spacing, removed trailing line ending, and inserted Content-Type MIME header.

related to multipart form data, and rejects requests that are not compliant. For those requests that are compliant, optional message body components and fields that must be ignored are removed. This ensures that problematic message constructs are removed before they can induce differential parsing behavior between the WAF and the web application framework. The `HTTP-Normalizer`'s functionality could be expanded for each content-type to provide protection against other popular content types as well. We are introducing a proof of concept description in this section, and use `multipart/form-data` as demonstration.

7.1 Overview of HTTP Normalizer

The HTTP multipart body normalizer serves two primary functions: normalizing multipart request bodies, and rejecting non-compliant requests. Upon receiving an HTTP request, the normalizer separates the request's headers and body using the AIOHTTP Python library. It then parses and validates the request's `Content-Type` header, and if it indicates that the message body uses the MIME multipart encoding, it performs one of the following actions:

Normalization. The normalizer begins by parsing requests using a strict multipart MIME parser generated from the ABNF in the RFCs. The output data structure of the parser is capable of representing only the necessary components of a multipart message body. Deprecated and optional portions of the multipart message are not representable in this structure. The structure is then reserialized and forwarded to its destination. Crucially, the normalized request is deserialized from a data structure in which invalid state is not representable, so the output of the normalizer is never malformed.

Rejection. If the normalization process cannot be completed because necessary components of the request body uses a deprecated RFC feature or is malformed, the request is rejected.

Listing 5 demonstrates a sample malformed request before being passed to the normalizer. Listing 6 demonstrates how the sample request would look like after being normalized.

7.2 Normalizing Request Body

Currently, the normalizer applies only to requests that use a multipart transfer encoding. Future enhancements will extend this normalization capability to other content-types, aiming to demonstrate the viability of this approach in enhancing WAF security by rejecting malicious HTTP requests and reducing false positives through normalization. This methodology can be uniformly applied to all content types to achieve comprehensive normalization.

Table 6: Evaluation Metrics for HTTP Normalizer.

	Bypassable Requests
Normalized	8
Rejected	55
Total Attempts	63

7.3 Evaluation of Normalizer

We evaluated the normalizer first by verifying that it does not reject any of a set of valid multipart message bodies. We obtained these from the test suites of the Tornado web server, RStudio, and the multiparty multipart message parser.

We sampled bypasses from all 12 discovered bypass classes that we mentioned in subsection 6.3.1 and took 63 bypassable requests in total to test against the normalizer, and recorded whether the request was accepted, normalized, or rejected. Note that while we could mutate HTTP requests using our fuzzer and send all of them to the normalizer, it was unnecessary because: 1) even if a request bypassed the normalizer, it would have been blocked by the WAF since it was not among the found bypasses, and 2) the normalizer will ultimately be merged into existing WAF rulesets and does not need to be fuzzed separately because it is not a WAF itself, but a ruleset enforcer.

Table 6 demonstrates these evaluation metrics. Out of the 8 bypassable requests that are not rejected by the normalizer, all are blocked by Cloudflare's WAF, resulting in a 100% success rate.

7.4 Performance Cost

The `HTTP-Normalizer` was developed to demonstrate that all identified bypass techniques can be prevented if WAFs adhere to proper parsing standards. Our findings, including the lack of successful bypasses against AWS WAF, validate this hypothesis. The intent of the project was not to create a fully optimized, production-grade tool with minimal overhead and broad content-type support, but rather to prove that an effective solution is feasible and implementable for a specific content type. By strictly adhering to RFC grammar rules, the HTTP-Normalizer effectively eliminates all bypass methods. While it currently functions as an additional evaluation layer, its

rules can be seamlessly integrated into existing WAF rulesets with minimal performance impact.

8 Discussion

In this section, we discuss key considerations and challenges identified in our work, as well as exploring potential for future research and development based on our findings.

8.1 Potential for WAF and Framework Fingerprinting

Previous research has shown that it is possible to use HTTP parsing discrepancies to fingerprint web servers by utilizing a combination of requests known to trigger specific actions on target web servers [22]. The findings of our work also show that parsing discrepancies are present in WAFs and web application frameworks, which opens up the possibility of fingerprinting them by sending requests known to bypass and those that do not bypass WAFs. A concrete example from our findings would be sending a request to the target from the JSON bypass class under the Field Name Manipulation category, which we know will only bypass and succeed if the target runs on Google Cloud Armor with the Spring Boot framework. We can send the request and then analyze the response to identify whether the underlying WAF is Google Cloud Armor. Future work could expand on this and propose a fingerprinting methodology on WAFs by using the findings of our work.

8.2 Practical Considerations

We focused on prominent web frameworks and their default or most popular parser packages. While multipart forms are common, not all frameworks have dedicated parsers for them, often leading to custom implementations from developers. These custom parsers can increase the risk of developer errors, potentially allowing bypassed requests to be parsed successfully server-side. However, our work did not examine custom parsers, leaving their implications outside the scope of our findings.

8.3 Limitations

This work primarily focuses on HTTP/1.1 web requests, thus future research could investigate HTTP/2. Also, the scope of WAFs tested may not cover all available systems. Future work could expand the range of tested frameworks and target additional content-types beyond those examined in this work. Exploring these areas would provide a more thorough understanding of WAF vulnerabilities and the potential for bypasses.

8.4 Usability vs. Security Debate

While our findings show that WAFs can prevent these bypasses by adhering to proper RFC standards, real-world implementations may face challenges due to compatibility constraints or customer-specific requirements. Therefore, we do not blame WAF vendors for these bypasses, as they often face trade-offs between security and operational needs. However, the our proposed defense mechanism, the `HTTP-Normalizer`, demonstrates that all identified bypasses in this work are theoretically defendable, ensuring enhanced security without significantly impacting usability. Also, note that not all WAFs are susceptible to every type of bypass, suggesting that a well-designed parser adhering to all relevant RFC rules can mitigate these vulnerabilities.

9 Ethics Considerations and Disclosure Information

All attacks and bypasses in this paper were conducted in a controlled environment using our infrastructure, ensuring no external impact. For our analysis, no malicious payloads were sent to the tested websites, and each site received no more than 10 requests.

All WAF tests adhered to bug bounty protocols, and bypasses were disclosed to the affected vendors, who subsequently acknowledged the existence of these vulnerabilities.

- **Google Cloud Armor** classified our report as a Tier 1, Priority 1, Severity 1 vulnerability under the "Insecure by Default" category and rewarded us with a bug bounty.
- **ModSecurity** acknowledged all bypasses we reported in CRS 3.3.
- **Cloudflare** confirmed the reported bypasses and stated that they are working on a fix.
- **Microsoft Azure** acknowledged the bypasses in CRS 3.0, which was the default ruleset at the time of our study. However, they are retiring CRS 3.0 and transitioning to DRS 2.1, an enhanced ruleset based on CRS 3.2, which addresses these vulnerabilities.

Our results and source code, including fuzzer input grammars, are publicly available[4]. Sensitive bypass requests will remain restricted until the vulnerabilities are resolved.

10 Conclusion

This work has highlighted the significant impact of parsing discrepancies on the effectiveness of Web Application Firewalls (WAFs) in protecting against cyber threats. Our experiments across major WAFs, particularly with `multipart/form-data`, `application/json`, and `application/xml` content-types, resulted in 1207 bypasses. These bypasses were successfully parsed by our target web application frameworks employing default or popular request parsers. This finding indicates a critical vulnerability in WAFs' ability to uniformly interpret and filter web requests.

This paper's investigation into popular WAFs, including Cloudflare, Cloud Armor, AWS WAF, Azure WAF, and ModSecurity, highlights significant concerns regarding parsing discrepancies in their request analysis. These discrepancies pose a substantial risk, as WAF users may falsely believe they are protected against common attack payloads.

As attackers employ increasingly sophisticated methods, the development of dynamic, intelligent WAFs becomes important. Our research not only contributes to the understanding of current WAF limitations, but also proposes effective solutions to defend against the identified threats. This work emphasizes the critical need for improved parsing consistency in WAFs to ensure robust protection against these attacks. Our proposed `HTTP-Normalizer` offers a promising solution, balancing security and usability by enforcing strict compliance with RFC standards.

[4] https://github.com/sa-akhavani/waffled/

Acknowledgment

This work was supported by the National Science Foundation under grant 2329540.

References

[1] 2021. A Forgotten Invisibility Cloak. https://soroush.me/downloadable/A_Forgotten_HTTP_Invisibility_Cloak_v1.1.pdf [Online; accessed 16. Nov. 2023].
[2] 2022. Google WAF bypassed via oversized POST requests. https://portswigger.net/daily-swig/google-waf-bypassed-via-oversized-post-requests
[3] 2024. OWASP CRS | OWASP Foundation. https://owasp.org/www-project-modsecurity-core-rule-set/ [Online; accessed 3. Sep. 2024].
[4] Tim Bray. 2017. *The JavaScript Object Notation (JSON) Data Interchange Format.* Request for Comments RFC 8259. Internet Engineering Task Force. doi:10.17487/RFC8259 Num Pages: 16.
[5] D. Crocker and P. Overell. 2008. Augmented BNF for Syntax Specifications: ABNF. Request for Comments: 5234. https://www.rfc-editor.org/rfc/rfc5234 Accessed: 2025-01-07.
[6] Ned Freed and Dr. Nathaniel S. Borenstein. 1996. *Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types.* Request for Comments RFC 2046. Internet Engineering Task Force. doi:10.17487/RFC2046 Num Pages: 44.
[7] Ned Freed and Nathaniel S. Borenstein. 1996. *Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies.* Request for Comments RFC 2045. Internet Engineering Task Force. doi:10.17487/RFC2045 Num Pages: 31.
[8] Mattias Grenfeldt, Asta Olofsson, Viktor Engström, and Robert Lagerström. 2021. Attacking websites using HTTP request smuggling: empirical testing of servers and proxies. In *2021 IEEE 25th International Enterprise Distributed Object Computing Conference (EDOC)*. IEEE, 173–181.
[9] Mark Handley, Vern Paxson, and Christian Kreibich. 2001. Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics. In *USENIX Security Symposium*.
[10] Internet Engineering Task Force (IETF). 2025. RFC Process. https://www.ietf.org/process/rfcs/ Accessed: 2025-01-07.
[11] Bahruz Jabiyev, Steven Sprecher, Kaan Onarlioglu, and Engin Kirda. 2021. T-Reqs: HTTP Request Smuggling with Differential Fuzzing. In *Conference on Computer and Communications Security (ACM CCS)*.
[12] Suman Jana and Vitaly Shmatikov. 2012. Abusing File Processing in Malware Detectors for Fun and Profit. In *IEEE Symposium on Security and Privacy*.
[13] James Kettle. 2020. HTTP Desync Attacks: Request Smuggling Reborn. https://portswigger.net/research/http-desync-attacks-request-smuggling-reborn. Accessed: 2025-01-07.
[14] Amit Klein. 2020. HTTP Request Smuggling in 2020: New Variants, New Defenses, and New Challenges. In *Black Hat USA*. https://www.blackhat.com/us-20/briefings/schedule/#http-request-smuggling-in-2020--new-variants-new-defenses-and-new-challenges-20019 Accessed: 2025-01-07.
[15] Amit Klein. 2020. HTTP Request Smuggling in 2020: New Variants, New Defenses, and New Challenges. https://i.blackhat.com/USA-20/Wednesday/us-20-Klein-HTTP-Request-Smuggling-In-2020-New-Variants-New-Defenses-And-New-Challenges-wp.pdf. Accessed: 2025-01-07.
[16] Ed Levinson. 1998. *The MIME Multipart/Related Content-type.* Request for Comments RFC 2387. Internet Engineering Task Force. doi:10.17487/RFC2387 Num Pages: 10.
[17] Larry M. Masinter. 2015. *Returning Values from Forms: multipart/form-data.* Request for Comments RFC 7578. Internet Engineering Task Force. doi:10.17487/RFC7578 Num Pages: 15.
[18] Keith Moore. 1996. *MIME (Multipurpose Internet Mail Extensions) Part Three: Message Header Extensions for Non-ASCII Text.* Request for Comments RFC 2047. Internet Engineering Task Force. doi:10.17487/RFC2047 Num Pages: 15.
[19] Hoai Viet Nguyen, Luigi Lo Iacono, and Hannes Federrath. 2019. Your Cache Has Fallen: Cache-Poisoned Denial-of-Service Attack. In *ACM Conference on Computer and Communications Security*.
[20] Zhenqing Qu, Xiang Ling, and Chunming Wu. [n. d.]. AutoSpear : Towards Automatically Bypassing and Inspecting Web Application Firewalls. ([n. d.]). https://i.blackhat.com/Asia-22/Thursday-Materials/AS-22-Qu-AutoSpear-Towards-Automatically-Bypassing-and-Inspecting-Web-Application-Firewalls.pdf
[21] Henry S. Thompson and Chris Lilley. 2014. *XML Media Types.* Request for Comments RFC 7303. Internet Engineering Task Force. doi:10.17487/RFC7303 Num Pages: 35.
[22] Cem Topcuoglu, Kaan Onarlioglu, Bahruz Jabiyev, and Engin Kirda. 2024. Untangle: Multi-Layer Web Server Fingerprinting. In *Network and Distributed Systems Security Symposium (NDSS)*.
[23] Qi Wang, Jianjun Chen, Zheyu Jiang, Run Guo, Ximeng Liu, Chao Zhang, and Haixin Duan. 2024. Break the Wall from Bottom: Automated Discovery of Protocol-Level Evasion Vulnerabilities in Web Application Firewalls. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 129–129.
[24] Zhongjie Wang, Shitong Zhu, Yue Cao, Zhiyun Qian, Chengyu Song, Srikanth Krishnamurthy, Kevin Chan, and Tracy Braun. 2020. SymTCP: Eluding Stateful Deep Packet Inspection with Automated Discrepancy Discovery. In *Network and Distributed Systems Security Symposium (NDSS)*.

A Tested Parsers

Table 7: List of tested parsers, their weekly downloads as of July 2024 collected from NPM and PyPI Stats [5]website, and whether they are vulnerable to at least one bypass in our study or not.

Parser	Version	Vulnerable	Downloads
Node.js			
Busboy	1.6.0	Yes	8,639,607
Formidable	3.5.1	Yes	6,998,735
fast-xml-parser	4.3.3	Yes	14,254,527
Python			
xmltodict	0.13.0	Yes	10,374,671
xmlminidom	3.10.12	Yes	n/a
python-multipart	0.0.9	Yes	4,288,034

B Web Framework Parsing Methods

Table 8 summarizes the methods used for parsing request bodies in our work in each framework. If a content-type is not supported by the default parser of a framework, popular third-party parsers are used instead.

[5]https://www.npmjs.com and https://pypistats.org

Table 8: Request Body Parsing Methods in Tested Frameworks

Framework	Content-Type	Parsing Method
Express (Node.js)	application/json	`request.body`
FastAPI (Python)	application/json	`request.json()`
Laravel (PHP)	application/json	`$request->collect()`
Spring Boot (Java)	application/json	`Defined a class and used @RequestBody`
Laravel (PHP)	multipart/form-data	`$request->collect()`
Express (Busboy)	multipart/form-data	`form.parse(request)`
FastAPI (Python)	multipart/form-data	`field1: str = Form("none")`
Flask (Python)	multipart/form-data	`request.form`
Spring Boot (Java)	multipart/form-data	`@RequestParam(value = "field1") String field1`
Laravel (PHP)	application/xml	`$request->getContent()`
Spring Boot (Java)	application/xml	`Defined a class and used @RequestBody`
Flask (xmltodict)	application/xml	`parse(request.data)`
Gin (Go)	application/xml	Defined an `xmlform` with `field1` and `field2`