Monitoring and Analyzing SSDP DDoS Amplification Attacks

An Empirical Study of Reflective Amplification Traffic Using Honeypots

Tim Guldenmundt[1]

Supervisor(s): Harm Griffioen[1]

[1]**EEMCS, Delft University of Technology, The Netherlands**

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 21, 2026

Name of the student: Tim Guldenmundt
Final project course: CSE3000 Research Project
Thesis committee: Harm Griffioen, Mitchell Olsthoorn

Abstract

Distributed Denial-of-Service amplification attacks are a growing threat to modern-day network infrastructure, with the Simple Service Discovery Protocol (SSDP) being one of the most abused protocols for these attacks. This paper revisits the SSDP amplification attacking landscape by designing, implementing and deploying a honeypot system, to capture real-world attack traffic.

The collected data is analyzed to study how adversaries select their amplifiers, what techniques they use and what countries and industry sectors are most often targeted in these attacks. The results show a clear preference for high amplification reflectors, an increase in the usage of subnet-wide "carpet bombing" strategies, and a concentration of attacks on ISP and hosting infrastructures.

Altogether this research presents an updated overview of the modern-day SSDP DDoS amplification attack landscape.

1 Introduction

Distributed Denial-of-Service (DDoS) attacks are one of the most damaging threats to modern network infrastructure. Large attacks can take down significant parts of the internet. Both the scale and frequency of these attacks have been growing in recent years. In late 2025 Cloudflare, one of the largest network security platforms, reported that the total number of DDoS attacks had doubled compared to 2024 [4]. Most of these attacks make use of so-called amplification attacks, where third party services are abused to amplify the total volume of attack traffic.

Many different protocols are known to be abused as an amplification factor for these types of attacks, with the Simple Service Discovery Protocol (SSDP) being one of the most common. According to a study done by Microsoft in 2022 about 15% of all amplification attacks are making use of SSDP [2].

Previous studies have looked at the landscape of these amplification attacks by deploying amplification DDoS honeypots [7; 8; 10; 12]. However most of these studies have been done multiple years ago, with the most recent studies dating back to 2021, and with the fast rise in frequency of DDoS attacks, it is important to take a new look at the methods adversaries use in their attacks.

This study thus aims to revisit the DDoS amplification landscape and will specifically look at the use of SSDP in amplification attacks. It examines the methods and strategies used by adversaries in real-world DDoS attacks, monitors changes compared to previous studies and tries to identify new, possibly previously overlooked patterns. This leads to the research question of this study: *How is the Simple Service Discovery Protocol abused for DDoS attacks in practice by adversaries?*

To answer this question this study will try to find answers to the following sub questions:

- Are adversaries actively selecting amplifiers with a high amplification rate for their attacks?
- Do attack duration and request rate differ between attacks?
- What strategies do adversaries use when executing SSDP amplification attacks?
- Who are the typical victims of DDoS attacks using SSDP amplification? Are specific industries targeted more often?

To answer these questions, this research first explains the relevant background on DDoS attacks and honeypot systems. It then presents the design of the honeypot system that was deployed and used to collect the data for this research in section 3. The fourth section describes the data collection process, the attack classification algorithm used to distinguish between attacks and scanners, and the results from the collected data. Section 5 discusses the responsible research aspects of this study, such as the contribution to attacks, privacy concerns and the reproducibility of the research. The next section interprets the results and places them in a broader context, including the results of previous research. Finally, the last section summarizes the main conclusions from this study, outlines it limitations and makes suggestions for future research.

2 Background

Denial-of-service (DoS) attacks aim to disrupt the legitimate use of a service. A classic example of a DoS attack would be flooding a network with traffic, thereby exhausting its bandwidth and preventing legitimate traffic from being served. Distributed Denial-of-Service (DDoS) attacks make use of multiple attacking entities. Attackers first exploit security vulnerabilities on machines and infect them with the malicious code. These new agent machines are then used to send the attack packets towards the victims [14].

In order to generate the biggest DDoS attack possible, attackers often make use of so-called Distributed Reflective Denial-of-Service (DRDoS) attacks. DRDoS attacks (also known as amplification attacks) make use of the fact that some public servers respond to specific UDP network protocol requests without validating the identity of the sender [18]. Attackers can misuse this attribute by sending requests while spoofing the victim's IP. The public server will then unknowingly send their response to the victim instead of the attacker. The main benefit for attackers is that some of these servers respond with larger messages than the original request thus amplifying their attack size. This paper will look specifically at the Simple Service Discovery Protocol (SSDP), which is one of the more frequently used protocols for amplification attacks.

SSDP is a UDP-based discovery protocol used by devices to automatically discover each other on a network. When a device is searching for a specific service, it sends a multicast 'M-SEARCH' request on port 1900. Any listening device that is capable of providing that service then responds with a unicast response directly to the requesting device [17]. While the protocol was designed to only work on local networks,

devices will also respond to unicast 'M-SEARCH' requests from outside their local network. This is what allows attackers to misuse SSDP as a reflector for amplification attacks.

To collect real-world attack data, an SSDP honeypot system was designed and implemented in this research. A DDoS honeypot is a system designed to emulate a vulnerable amplification service to attract and log attack traffic. An amplification honeypot should have realistic amplifying behaviour to be attractive to attackers. However, exposing amplifying services to the internet will aid attackers in executing their attacks. To minimize the harm done by deploying a honeypot, strict rate limiting is often applied [15].

3 Honeypot Design

The honeypot system is made in Go [6] and is designed to emulate multiple devices responding to an 'M-SEARCH' request. By exposing this system on multiple accessible IP addresses, scanning traffic from attackers looking for these services and reflective traffic from actual DDoS attacks can be collected. To increase the visibility of the honeypot, it is deployed with 17 public IP addresses; this makes it easier for attackers to find the system. An overview of the system design can be seen in Figure 1.

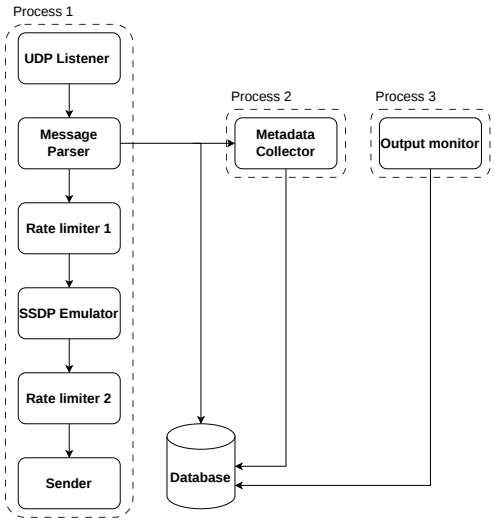

Figure 1: A system design diagram of the honeypot system deployed in this research

3.1 Amplification factor

To measure the amplification factor, a metric named the bandwidth amplification factor (BAF) [18] is used:

$$BAF = \frac{len(response)}{len(request)} \quad (1)$$

The honeypot IP's are split into 2 groups: small or large response. To be specific, the small response group will emulate a response of 3 devices, and the large group emulates 10 devices. Table 1 shows the response size in Bytes and the BAF for a typical request with a size of 94 Bytes.

Group	Response Size (Bytes)	BAF
1	1329	14.14
2	4616	49.11

Table 1: Response size and BAF of the two IP groups for a typical request of 94 bytes.

3.2 Rate limiting

To minimize any harm caused by this research, two rate limiting mechanisms are used:

1. The first rate limiter tracks the amount of requests per **(Source IP, Destination IP)** pair. When a pair generates more than 100 request in a single minute, the pair is blocked for 1 hour. With a destination IP from the large response group, this would block the pair after reaching a bandwidth per minute of:

   $$100 * 4616B = 0.4616MB$$

   Besides the source IP, its **/24** subnet is also blocked for 1 hour. This aims to prevent the possibility of carpet bombing attacks, that target multiple IP's within a subnet instead of a single IP [8], avoiding the rate limit system.

2. The second rate limiter monitors the output of all honeypot IP's and limits the total output to 1 MB/s. This is primarily used to minimize the possibility of an attacker using all the honeypot IP's in a single attack, but also is used as a safety guard in case the first rate limiter fails.

3.3 Logging

All incoming requests to the honeypot are logged to a ClickHouse [3] database. For each request the system logs the: **Source IP, Source port, Destination IP, Time, Size** and **Request Type**. Furthermore for each source IP, the results of a reverse DNS lookup are stored, alongside the Geo location and Autonomous System Number (ASN), which are collected from the Maxmind GeoLite2 [13] database.

ClickHouse was chosen because it is a column-oriented analytical database that has been optimized for very high write throughput. A honeypot system can expect to receive high amounts of traffic during an attack and the system must therefore be able to store large amounts of requests without dropping data.

To decouple the traffic collection from the data storage, a NATS stream [19] was deployed between the main honeypot process (seen as Process 1 in Figure 1) and the database. NATS is a lightweight publish-subscribe messaging system allowing the honeypot to asynchronously send messages to the stream, without directly inserting them in the database. A separate consumer service listens to the stream and writes the request in batches to the database.

3.4 Deployment

The honeypot system was deployed on the 29th of May 2026, on a machine running Ubuntu 24.04.4. The system has 17 IP address on which it listens for UDP messages on port 1900, 9 of these have been assigned to the small response group and 8 to the large response group, all of these are registered in the Netherlands.

4 Attack Classification

In order to analyze the collected data, a classification algorithm is needed to identify and group requests that belong to the same attack. The algorithm for this research works in two stages.

First the classification algorithm will use a combination of **(Source IP prefix, Destination IP)** to group attacks. The prefix of the source IP is used instead of the full address to avoid splitting carpet bombing attacks into separate attacks. The destination IP is included to correctly identify attacks over scanning traffic. Without including the destination IP scanners that scan all honeypot IP's in quick succession would be indistinguishable from small attacks.

To decide which requests belong to the same attack, two parameters need to be defined.

1. τ: the maximum allowed time gap in seconds between two consecutive requests within a session.
2. N_{min}: the minimum amount of requests for a sessions to be considered as an attack

To determine suitable values for these parameters, a parameter sweep will be performed for both parameters over a range of possible values. First a sweep is performed for τ, and the value is chosen at the point where the number of detected sessions no longer decreases significantly. At this point most fragmented sessions have been merged, and choosing a higher value increases the risk of merging two separate attacks. Next the parameter sweep is repeated for N_{min} with τ already set. Again the value is chosen to a point where the session count stabilizes, indicating that most scanning traffic has been excluded. Increasing its value risks excluding small attacks from the data.

The second stage of the classification algorithm merges sessions that were classified in the first stage across destination IP's. Since the system is deployed with 17 IP addresses, some attacks might use multiple honeypot IP's in a single attack, and have thus been classified as separate attacks in the first stage. Two sessions are merged if they share the same source IP prefix, and their time frames overlap or are within τ of each other.

5 Experimental Setup and Results

The honeypot system started collecting data on May 29, 2026, and ran for 17 days until June 16, 2026. Due to a crash on June 10, 2026 the system stopped collecting data for approximately 6 hours. In total the honeypot collected 84,532,921 requests during the measurement period.

5.1 Attack classification parameter sweeps

To determine the correct values for the attack classification algorithm's parameters, a sweep was performed for both τ and N_{min} on the collected data as described in section 4. Figure 2 shows the results for the parameter sweep of τ. Based on the parameter sweep, the value of τ is set at 120.

Figure 3 shows the sweep performed for the N_{min} parameter, based on τ set at 120 seconds. The N_{min} parameter is set at 5 based on the results from the sweep.

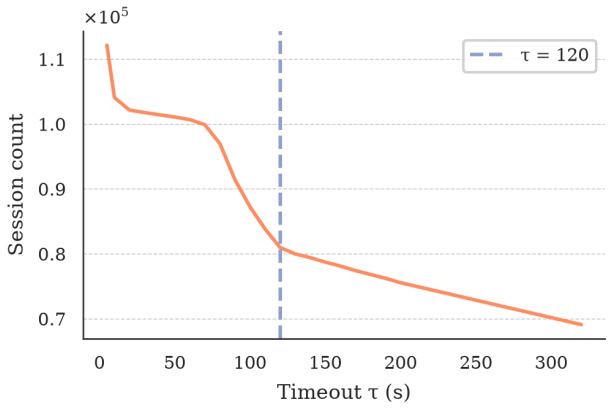

Figure 2: Detected session counts for different value's of τ.

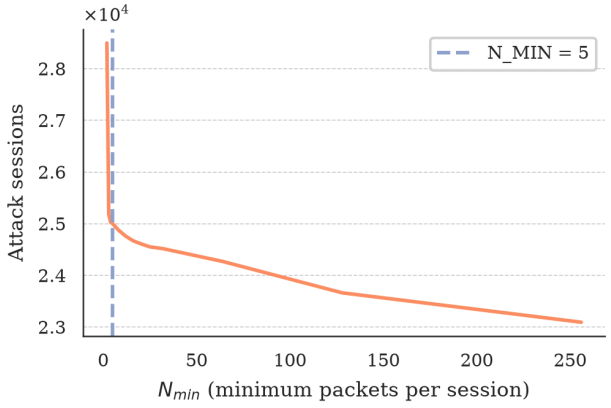

Figure 3: Detected attack sessions for different value's of N_{min}.

5.2 Results

Using the classification algorithm described in section 4, all requests can be classified as either attack or scanning traffic. Figure 4 shows the scanning activity for the duration of the system deployment by plotting the unique IP addresses per 6 hours that are not classified as an attack. Note that the drop to 0 at 2026-06-10 is caused by the brief system downtime. Figure 5 shows the amount of attack session during the deployment.

Amplification factor usage

To identify if adversaries have a preference for higher amplification reflectors, the systems IP addresses were split into two groups: a high and low bandwidth amplification factor group. Figure 6 shows a comparison between the usage of the two groups in terms of total attack sessions and total received attack packets. Note that the low-BAF group consisted of 1 more IP address, and thus had slightly better visibility for attackers. Both Figure 6a and Figure 6b show a higher usage of the high amplification group. However there is a much larger difference in total attack packets between the groups than total attack sessions.

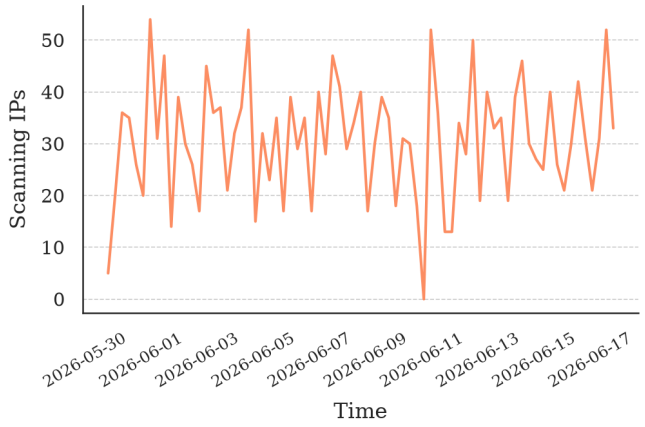

Figure 4: Unique non-attack IP's per 6 hour interval

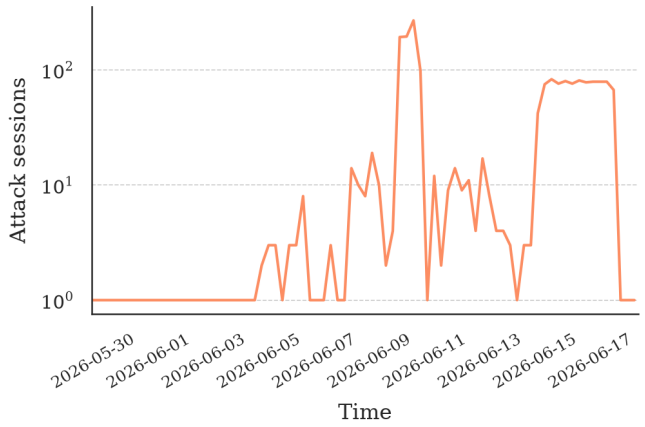

Figure 5: Detected attack sessions per 6 hour interval

Attack session characteristics

Table 2 shows the statistics for three attack sessions characteristics: session duration, total session requests and request rate.

Table 2: Distribution of attack session duration, request count and request rate.

Statistic	Duration (s)	Requests	Average Rate (req/s)
Mean	208.3	45 644.1	567.9
Median	123.0	13 260.0	89.5
Max	18 817	3 472 065	70 000

Across all three characteristics the mean is notably higher than the median, indicating that there are a few extreme outliers pulling the mean upwards. Figure 7 shows the distribution of the session duration. The volume of attacks spikes slightly around 30 s, 60 s, 120 s, and increases significantly around the 300 s (5 min) mark.

Attack strategy

Table 3 shows the amount of detected attacks targeting more than one victim IP address. The majority of attacks (77.0%) target more than one IP address and belong to so-called carpet bombing attacks.

Table 3: Amount of attacks targeting either a single IP address, or multiple addresses within a /24 subnet

Attack type	Count	Percentage
Single-source (= 1)	426	23.0%
Multi-source (> 1)	1 426	77.0%
Total	1 852	100%

Table 4 summarizes the use of source port randomization across attacks. Most attacks (84.6%) use multiple source ports. This technique is often used to avoid simple packet based filtering techniques [18].

Across all 1 852 detected attacks, only 325 unique victims are observed. 145 IP prefixes have been victim to multiple attacks, with the most targeted IP falling victim to 31 distinct attacks with an average gap of around 6 hours between successive attacks.

Table 4: Amount of attacks using either a single source port, or multiple source ports

Attack type	Count	Percentage
Single-port (= 1)	286	15.4%
Multi-port (> 1)	1 566	84.6%
Total	1 852	100%

Victim analysis

To identify what countries are most common victim to SSDP DDoS attack, the MaxMind GeoLite2 [13] database is used to map victim IP's to their countries. Table 5 shows the top 5 most targeted countries by total requests from detected attack sessions. For the complete table see Appendix A.

Country	Request count	%
United States	28 826 736	34.10
Hong Kong	14 253 732	16.86
Japan	9 777 241	11.57
Singapore	7 780 831	9.20
Australia	4 071 367	4.82

Table 5: Top 5 victim countries by total SSDP requests from sessions classified as an attack and percentage of overall volume.

Besides the victim country, the industry sector is also analyzed. First the victim IP's are mapped to their Autonomous System Number (ASN) using the GeoLite2 database [13]. Then the Stanford ASDB database [1] is used to map the ASN to an industry sector. ASDB is a research dataset that maps ASN to their organization and three industry categories using a machine learning algorithm. Table 6 shows the most targeted sectors by total requests from sessions classified as attacks, observed by the honeypot system.

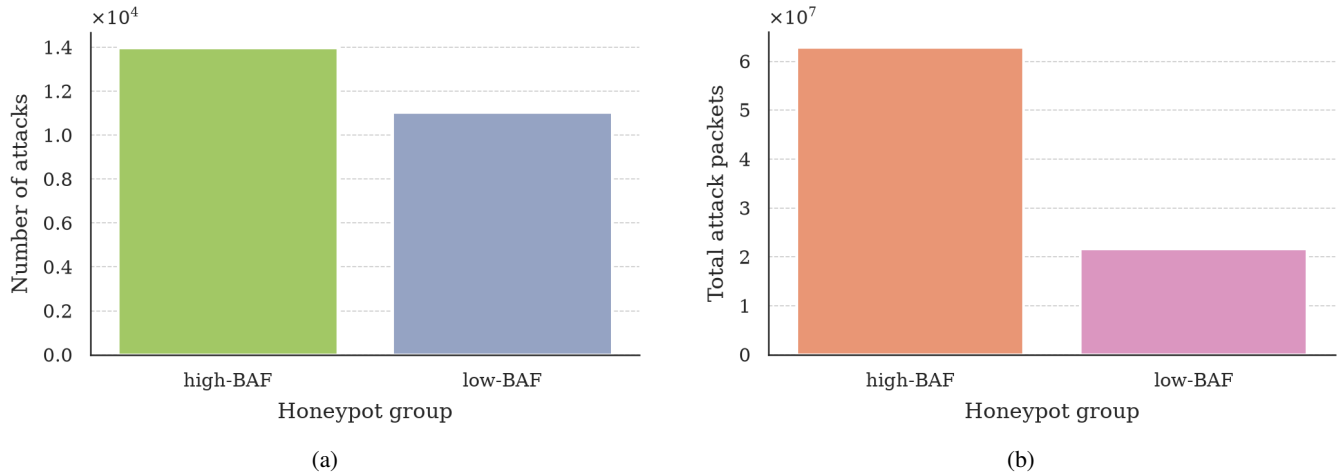

Figure 6: Comparison of high- and low-BAF honeypot group usage: (a) number of attack sessions per group and (b) total packets received from attacks per group.

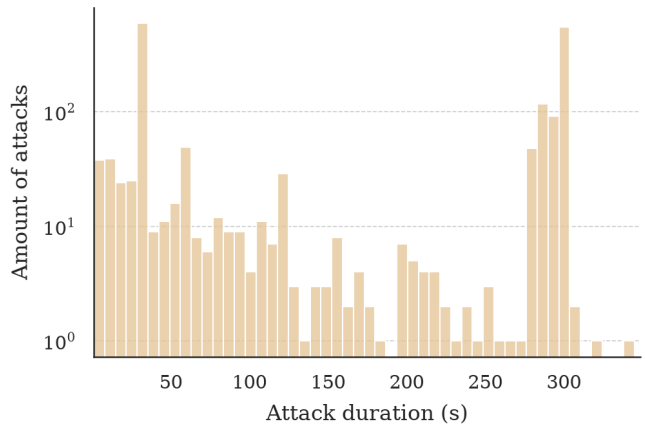

Figure 7: Distribution of attack session durations. Sessions above the 97th percentile (*>348.8*s) are excluded for visual clarity (*56* sessions excluded).

Sector	Request count	%
Other / unknown	48 683 758	57.59
ISP / consumer access	17 700 125	20.94
Hosting / cloud / IT	17 286 806	20.45
Academic / research	739 709	0.88
Government / public sector	100 257	0.12
Financial / banking	22 266	0.03

Table 6: Victim sectors by total SSDP requests from sessions classified as an attack and percentage of overall volume.

Most of the victim IP's (57.59%) could not be classified to a specific sector. These could be attacks to smaller organizations that are not included in the ASDB dataset. The rest of the victim IP's are almost completely targeted at ISP and consumer access networks (20.94%) and hosting, cloud, and IT infrastructure (20.45%).

6 Responsible Research

This section reflects on some of the ethical implications of deploying an amplification honeypot and collecting network measurement data. Furthermore it addresses the use of artificial intelligence during this research, and the reproducibility.

6.1 Attack contribution

By deploying a honeypot that responds to SSDP requests and participates in amplification attacks, this research could benefit attackers and cause harm to victims. To mitigate this risk, the honeypot system is deployed with strict rate limiting, as described in subsection 3.2. While the honeypot still partially contributes to the attack, the maximum bandwidth will never exceed 1 MB/s. With modern DDoS attacks now increasingly exceeding 1 TB/s [4], the contribution of the system to any single attack is almost negligible. However, it is important to acknowledge that, even though it is small, there is still unwanted traffic reflected towards victims.

6.2 Data collection and privacy

The honeypot records the source IP addresses for each incoming request. Under the General Data Protection Regulation (GDPR) [5] IP addresses are considered personal data since they can be used to identify individuals. However, the privacy concerns from this data are limited in practice. All collected IP addresses can be categorized as victim IP's and scanner IP's. The victim IP's are spoofed and belong to uninvolved third parties. The scanner IP's could belong to attackers scanning for amplifiers, but are equally likely to belong to legitimate network scanners. In both cases the IP's cannot be used to identify parties involved in malicious activity.

Access to the database containing the original IP addresses is limited, and any published data contains hashed IP's to anonymize the data.

6.3 Reproducibility

The complete honeypot source code is publicly available at: https://github.com/tgulden05/SSDP-Honeypot

The repository includes the configuration files used for this research, with the IP addresses and credentials left empty.

A limited version of the collected data is also available at: https://zenodo.org/records/20788351

All IP's in the public dataset have been hashed. This allows requests from the same source address to be identified without publishing the raw IP addresses.

While the experimental setup can be reproduced, the exact replication of the observed traffic is not possible, as incoming traffic depends on many outside factors. Deploying this system in a similar environment at a different time will possibly result in different traffic patterns.

6.4 Use of artificial intelligence

Some AI generated code has been used in the honeypot system. No AI was given direct write access to the repository: all generated code was first reviewed before being used, therefore full responsibility is taken for the behaviour of the deployed honeypot. No AI has been used for: creating the system design, analyzing collected data and drawing conclusions. AI has been used to check this paper for grammatical errors and suggest better sentence structures.

7 Discussion

Figure 6 shows that adversaries do tend to make more use of reflectors with a higher bandwidth amplification factor (BAF). This result is consistent with findings from Griffoen et al. [7], who found that attackers tend to look for higher amplification SSDP reflectors. Interestingly the difference in total attack packets is substantially bigger than the difference in total attacks between the two groups. This pattern could have multiple explanations:

1. Attackers do identify the high BAF reflectors and use them more intensively during their attacks, but do not exclude the lower BAF reflectors altogether.
2. Only the more sophisticated attackers, responsible for the largest attacks, make a selection based on amplification factor.

The attack session duration distribution (Figure 7) shows spikes around the 30 s, 60 s, 120 s and 300 s. These attacks could correspond to attacks from DDoS-for-hire services [11], which allow customers to launch attacks for a predefined duration depending on the purchased subscription. Alternatively, these values could appear more frequently because they correspond to convenient, round minute durations, which could naturally be selected more often by attackers.

The median attack duration measured by this research (123.0 s, seen in Table 2) is substantially lower than those reported by previous research. Heinrich et al. [8] measured a median duration of 612.5 s, and Jonker et al. [10] observed a median of 255 s. While this could suggest that the typical DDoS attacks have decreased in duration, the differences could also result from a difference in the attack classification algorithm.

The median attack request rate (89.5 req/s, seen in Table 2) indicates that most attackers are using large lists of amplifiers in their attack. Even with the response from the high-BAF honeypot addresses, this request rate would only result in an output of:

$$89.5req/s * 4616Bytes = 0.413MB/s$$

With the bandwidth capacities of current network infrastructures, this would only be a minuscule part of the bandwidth needed for a successful DDoS attack.

A significant finding of this research is the high percentage of attacks targeting more than one IP address within a **/24** subnet (77.0%, seen in Table 3). In 2021 Heinrich et al. [8] measured only 21.8% of attacks to be targeted at multiple IP addresses. This increase could signal a significant change in the attacking landscape: Where previously the carpet bombing attacks might have only been used by some of the more sophisticated attackers, nowadays it appears to be the standard for most DDoS attacks.

The largest share of attack traffic was directed towards IP addresses in the United States (34.1%, seen in Table 5), which is not surprising given that the United States also has the largest allocated IP address space [9]. Singapore and Australia appear notably often relative to their global network share. While not owning the largest global network space, both countries are among the highest in IPv4 addresses per capita. This could suggest that these countries are hosting large amounts of services, and therefore are more often victim to DDoS attacks.

Of the victim IP addresses that could be classified into an industry sector, the victims are almost evenly split between ISP and consumer access (20.94%, seen in Table 6) and hosting, cloud and IT infrastructure (20.45%). The high percentage of attack traffic aimed at ISP's is consistent with findings from Noroozian et al. [16] and could be due to the accessibility of DDoS-for-hire services, which are frequently used by individuals targeting other individuals in, for example, an online gaming context. The large share of traffic towards hosting and cloud infrastructure is most likely to reflect attacks against specific hosted services, such as websites or game servers.

8 Conclusions and Future Work

This research investigated how the Simple Service Discovery Protocol (SSDP) is abused by adversaries for DDoS amplification attacks in practice. To answer this question, a honeypot system was designed, implemented and deployed across 17 IP addresses. Over a 17 day measurement period the honeypot collected a total of 84 532 921 requests. These requests were then classified into either attack or scanning traffic using a classification algorithm.

8.1 Conclusions

Adversaries prefer high-amplification reflectors. The results show that adversaries have a clear preference for the higher amplification honeypot addresses, both in terms of total attack sessions and total attack packets. The difference in attack packets between the high and low amplification groups

was significantly larger than the difference between attack sessions.

Attack duration and request rate vary widely. The median attack duration and request rate both are significantly lower than their mean, indicating a small number of large outliers. There are several spikes in the session duration (30 s, 60 s 120 s and 300 s), which are consistent with fixed-duration attacks offered by DDoS-for-hire services. Furthermore the relatively low median request rate suggests that attackers are using large lists of amplifiers for their attack, instead of concentrating their requests towards a small selection.

Carpet bombing is becoming the default strategy for SSDP DDoS attacks. The amount of observed attacks targeting more than one IP within a **/24** subnet, also called carpet bombing attacks, is substantially higher than the findings of previous research. This indicates that the DDoS landscape is shifting towards carpet bombing being the standard approach for attacks using SSDP.

Victims are concentrated in a small section of countries and industry sectors. Within the 1,852 detected attacks, only 325 unique victims were observed, with the most targeted IP being victim to 31 distinct attacks. The United States fell victim to the most attack traffic, which is not surprising due to their large share of global network infrastructure. Among the victims that could be classified to an industry sector, almost all traffic was evenly split between ISP/consumer networks and hosting, cloud and IT infrastructure. Other sectors received very limited amounts of attack traffic.

8.2 Limitations

The results of this study are limited by several factors. First of all by only deploying the honeypot with 17 IP addresses, all from a single country, the visibility of the system for attackers is limited. While the amount of addresses appeared to be sufficient to capture some SSDP amplification attacks, it is very unlikely that it captured all ongoing attacks. A larger group of IP addresses would result in a more complete overview of the attacking landscape.

Second, the duration of the data collection period of approximately two weeks is relatively short compared to similar research. Some attack patterns might only become visible over a longer timescale, and the collected data from this research may represent a temporary trend instead of long term attacking behaviour.

A further limitation of the research is that part of the results depend on the values chosen for the parameters τ and N_{min}. There are multiple reasonable threshold values that would result in different splits or merges between attack sessions. A differently chosen value could therefore change some of the results from this research.

8.3 Contributions

This study contributes with an SSDP honeypot implementation that can be used to collect real-world data for studying SSDP amplification attacks. It further provides a simple attack classification method that groups requests based on source IP prefix and destination IP, and empirically selected parameters, that provides a separation between attack traffic and scanners. Finally this study offers an updated view of the DDoS amplification landscape, specifically for attacks abusing SSDP, and shows that the attack strategies of adversaries have shifted since the most recent surveys.

8.4 Future work

The results of this research leave several directions open for future research. First of all, the current classification algorithm is relatively simple, and choosing different values for the two parameters could result in different findings. Future research could take a closer look at the classification algorithm, identify its flaws and try to propose a better alternative.

Further more the honeypot deployed by this research only implemented SSDP emulation. Therefore no multi-protocol comparison can be made. Extending the implementation to include the emulation of other amplification protocols would allow for direct comparison between the behavior of adversaries across protocols.

With carpet bombing strategies showing a significant increase in popularity, future research could look into the effectiveness of this strategy against current detection and mitigation methods. Understanding why adversaries are shifting towards this attack strategy, and how it possibly evades some currently deployed defense systems, could benefit the development of new defenses.

Lastly, with this research only using a 17 day measurement period, future research could revisit these questions with a longer observation window. Some patterns may have remained hidden from this research due to the relatively limited amount of collected data, or some of the results may be influenced by extreme temporal outliers, instead of reflecting normal attacking behaviour.

A Victim country overview

Country	Request count	%
United States	28 826 736	34.10
Hong Kong	14 253 732	16.86
Japan	9 777 241	11.57
Singapore	7 780 831	9.20
Australia	4 071 367	4.82
United Kingdom	3 823 977	4.52
Germany	3 515 958	4.16
Seychelles	2 761 897	3.27
Macao	2 015 738	2.38
Czechia	1 876 124	2.22
United Arab Emirates	1 625 601	1.92
The Netherlands	1 357 810	1.61
Canada	531 717	0.63
Taiwan	372 366	0.44
Iran	247 101	0.29
Luxembourg	225 063	0.27
China	178 795	0.21
Poland	87 974	0.10
Vietnam	60 655	0.07
Italy	44 145	0.05
Portugal	43 819	0.05
Sweden	34 716	0.04
France	30 772	0.04
Hungary	24 720	0.03
Russia	23 026	0.03
Türkiye	18 903	0.02
Romania	14 244	0.02
San Marino	8 034	0.01
Pakistan	6 364	0.01
Estonia	5 918	0.01
South Korea	1 744	0.00
Brazil	1 559	0.00
Israel	954	0.00
Switzerland	454	0.00
Brunei	309	0.00
India	50	0.00

Table 7: Victim countries by total SSDP requests from sessions classified as an attack and percentage of overall volume.

References

[1] Stanford ASdb Dataset. https://asdb.stanford.edu/. Accessed: Jun. 18, 2026.

[2] Azure Network Security Team. Anatomy of a DDoS amplification attack. Microsoft Security Blog, may 2022. Accessed: 2026-06-21.

[3] ClickHouse, Inc. Clickhouse. https://clickhouse.com/, 2026.

[4] Cloudflare. 2025 q4 ddos threat report: A record-setting 31.4 tbps attack caps a year of massive ddos assaults. https://blog.cloudflare.com/ddos-threat-report-2025-q4/, 2025. Accessed: 2026-05-20.

[5] European Parliament and Council of the European Union. General data protection regulation (gdpr). https://gdpr-info.eu/, 2016. Regulation (EU) 2016/679. Accessed: June 2026.

[6] Google. The go programming language. https://go.dev/, 2026. Accessed: 2026-06-02.

[7] Harm Griffioen, Kris Oosthoek, Paul van der Knaap, and Christian Doerr. Scan, test, execute: Adversarial tactics in amplification ddos attacks. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, CCS '21, pages 940–954. ACM, 2021.

[8] Thiago Heinrich, Rafael R. Obelheiro, and Carlos A. Maziero. New kids on the drdos block: Characterizing multiprotocol and carpet bombing attacks. In *Information Systems Security and Privacy*, pages 284–302. Springer International Publishing, March 2021.

[9] Geoff Huston. IP addresses through 2025. APNIC Blog, January 2026. Accessed: 2026-06-20.

[10] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. Millions of targets under attack: A macroscopic characterization of the DoS ecosystem. In *Proceedings of the 2017 ACM Internet Measurement Conference*, IMC '17, pages 100–113. ACM, 2017.

[11] Mohammad Karami, Youngsam Park, and Damon McCoy. Stress testing the booters: Understanding and undermining the business of DDoS services. In *Proceedings of the 25th International Conference on World Wide Web*, WWW '16, pages 1033–1043. International World Wide Web Conferences Steering Committee, 2016.

[12] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. Amppot: Monitoring and defending against amplification ddos attacks. In *Research in Attacks, Intrusions, and Defenses (RAID 2015)*, volume 9404 of *Lecture Notes in Computer Science*, pages 615–636, Cham, 2015. Springer International Publishing.

[13] MaxMind. Geolite2 database. https://www.maxmind.com/en/home, 2026. Accessed: 2026-05-20.

[14] Jelena Mirkovic and Peter Reiher. A taxonomy of ddos attack and ddos defense mechanisms. In *ACM SIGCOMM Computer Communication Review*, pages 39–53, 2004.

[15] Marcin Nawrocki, John Kristoff, Raphael Hiesgen, Chris Kanich, Thomas C. Schmidt, and Matthias Wählisch. SoK: A data-driven view on methods to detect reflective amplification DDoS attacks using honeypots. In *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*, pages 576–591, July 2023.

[16] Arman Noroozian, Maciej Korczyński, Carlos Hernandez Gañan, Daisuke Makita, Katsunari Yoshioka, and Michel van Eeten. Who gets the boot? Analyzing victimization by DDoS-as-a-Service. In *Proceedings of*

the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID 2016), volume 9854 of *Lecture Notes in Computer Science*, pages 368–389. Springer, 2016.

[17] Open Connectivity Foundation. UPnP device architecture 2.0. https://openconnectivity.org/upnp-specs/UPnP-arch-DeviceArchitecture-v2.0-20200417.pdf, 2020. Accessed: 2026-06-01.

[18] Christian Rossow. Amplification hell: Revisiting network protocols for ddos abuse. In *Proceedings 2014 Network and Distributed System Security Symposium*. Internet Society, 2014.

[19] Synadia Communications, Inc. Nats: Cloud native messaging system. https://nats.io/, 2024. Accessed: 2026-06-21.