Measuring Global Risk of ICMP Amplification Attacks

Mostafa Dewidar
Stanford University

Ana Selvaraj
Stanford University

Victor Lin
Stanford University

Abstract

ICMP amplification is largely seen as an obsolete method of engineering Denial-of-Service (DoS) attacks. These attacks rely on a public server or amplifier that has not implemented valid mitigation strategies like restricting ingress ICMP traffic through a firewall. We sought to verify whether there is risk for such amplifiers in the IPv4 address space by measuring the Amplification Factor (AF) of public servers that respond to ICMP packets. Our experiments were concentrated on conducting the basic mechanism for the Smurf Attack against a set of 500,000 randomly chosen servers. For servers with high AFs, we determined their geographical location and ways they may be used as amplifiers in large-scale DoS attacks. The main contribution of our experiments was examining the minimal risk of ICMP amplifiers in the IPv4 address space using modern tools like Censys, Geolite and Zmap.

1 Introduction

DDoS (Distributed Denial-of-Service) is a widespread attack that tries to disrupt use of a network resource for legitimate users [13]. Famous incidents of DDoS attacks involve adversaries making use of the weak security of IoT devices to launch large-scale DDoS attacks like the Mirai botnet [7]. Another implementation of a DDoS attack that does not require the command-and-control style of a botnet is called an amplification attack. These attacks simply use reflectors/amplifiers that generate network traffic to a victim on behalf of them. These are now seen as a common type of attack, relying on IP spoofing and amplifiers (third-party entities that generate response packets to a query from an attacker that are larger than the original query) to reflect a large volume of packets to a victim. These attacks are generally more accessible than using botnets since they do not require infecting the amplifier.

The attack process of an amplification-based DoS attack involves scanning public IP addresses that interact with spoofed IP packets with large responses. An attacker would then use protocols that allow for amplification like DNS, NTP, etc. By consolidating the list of reflectors and protocols that generate the highest amplification volume, the attacker can estimate their approximate attack size. Efficient protocols for amplification-based DoS attacks have high Amplification Factors (AF), defined as the size of the response from the amplifier to the victim divided by the size of the query from the attacker to the amplifier.

A recent example of an amplification-based DoS attack is the "memcached" reflection attack [6]. This involved flooding a vulnerable open-source distributed memory caching system with UDP-based spoofed traffic. Its AF was measured to be around 10,000 to 51,000. In 2018, Akamai detected a 1.3 TBps DDoS attack which was said to be more than twice as large as the size of the Mirai botnet attack against Akamai in 2017.

Most amplified attacks like this rely on payload magnification of UDP-based protocols where the amplification relies on the larger packet size [13]. In contrast, TCP-based attacks are generally seen as difficult to implement well because of the three-way handshake.

1.0.1 ICMP Amplification Attacks

The first idea for the implementation of an amplification attack was through ICMP in the Smurf Attack [13]. This relies on flow multiplication where the number of packets is amplified rather than the payload size. The Smurf attack uses the ICMP echo broadcast where an echo request to a broadcast address with the spoofed IP of a victim would generate a large number of packets from the amplifier network to the victim.

This type of attack was popular in the late 90s and is seen as a "fixed" problem as router vendors implemented simple mitigation strategies. For example, Cisco

IOS version 12.0 and later disabled IP-directed broadcasts [17] by default. CDNs like Cloudflare also prevent Smurf attacks by preventing ICMP packets from reaching the targeted origin server with their firewall [10].

A limitation of this attack is that the attacker's network needs to allow IP spoofing, but most ISPs still allow it at the moment [16].

2 Related Work

Moon et al. [16] developed a framework, AmpMap, for measuring the risk of amplification that accounts for query and server variability. It randomly samples query patterns on a single server and then probes the queries with high AFs on other servers in a set to empirically quantify the risk of amplification attacks with a low network footprint. Their real-world measurements revealed how classic mitigation strategies like response rate limiting and secure configuration/setups can still leave risk for new amplification patterns. However, they only scanned thousands of servers for 6 UDP-based protocols.

In 2007, Kumar's work shows how accessible effective Smurf attacks are [14]. Through amplification factor calculations, they show how an attacker with a dial-up modem could overwhelm a class C network with a T1 or fractional T3 network link. They also discuss how smurf attacks rely on three components: the attack network, the intermediary network and the victim computer's network. Mitigation measures should be properly implemented at all these levels to mitigate Smurf attacks. The attacker's network firewall could inspect egress traffic to prevent spoofed traffic. The intermediary networks should not amplify the attack traffic while the final victim's network should filter out IP-directed broadcast addressed packets. However, in the real world, there is still risk for attacks because of ineffective prevention rules or misconfigured setup deployed at routers, firewalls or the IPS in the intermediary networks that the attack traffic passes through.

Bouyeddou et al. [9] developed an approach for detecting Smurf attacks through a Kullback-Leibler-based scheme. Rivas et al. [18] evaluate how CentOS performs under IoT based DDoS attacks. The attack scenarios were TCP-SYN flooding, UDP flooding, Ping flooding and Smurf Attack. They used different classes of addressing in their evaluation and conclude that CentOS 15 is a promising OS to host network services. Similarly, Gunnam et al. [12] investigate the performance of Microsoft's Windows Server 2012 R2 and Apple's Mac Server LION 10.7.5 under ICMP-based DDoS attacks. They found that both drop legitimate connections against Smurf attack traffic of above 150 Mbps and both types of server software need to implement more efficient protection mechanisms without relying on external security devices.

Bock et al. [8] scanned the entire IPv4 address space to find TCP non-compliant middleboxes that can act as mega-amplifiers for TCP-based amplification attacks. They found over 200 million IPv4 addresses (that are mostly middleboxes) to be such effective amplifiers that it's theoretically possible for them to have an infinite amplification factor. They discuss how middleboxes are an untapped threat where nation-states' censorship infrastructure can be redirected to reduce the security of the entire Internet.

3 Methodology

3.1 Setup

We used a VM to generate a list of 468 million IPv4 addresses that respond to ICMP echo requests using Zmap. It's important to note that these servers do not all respond with correct echo responses. They may respond with only failure messages but this does not prevent us from measuring their amplification factor. Our VM was in Stanford's network so all ICMP packets that were sent/received passed through its network.

Our experiments were all on a set of 500,000 addresses randomly drawn from the list of 468 million IP addresses. We reduced the address space so that we would be able to thoroughly profile these addresses.

By mapping each address to their geographical location using GeoLite2, we generated figures to compare the geographical distribution of our random set to the potentially amplifying subset of networks.

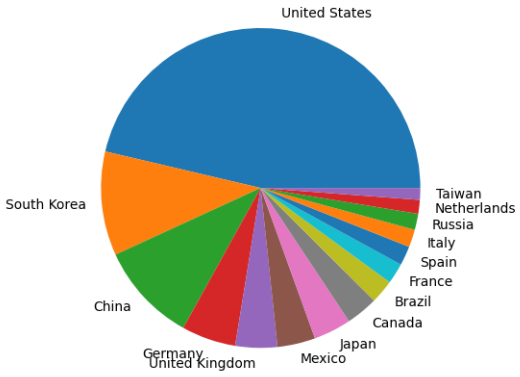

Figure 1: Geographical Distribution of our set of random 500,000 addresses

3.2 ICMP Request Implementation

We used multiprocess-based parallelism in Python to efficiently send echo requests from our VM to each IP address in the 500,000 set. We sent a single echo request to each IP in the list. To receive responses, we used the tshark utility to capture packets to and from our machine and create a capture log of all the packets. We used another Python script to analyze our log to measure amplification factors. The goal was to obtain amplifiers with AF > 1. This measurement excludes packets with ICMP type 3, which means the destination was unreachable, to avoid false positives [4].

3.3 Smurf Attack Implementation

The Smurf Attack is an amplification DoS attack pattern on ICMP that takes advantage of the broadcast addresses' ability to amplify packets [15]. Generally, an IP address that is part of a network will have an associated broadcast address that acts as an entry point into the network that forwards incoming packets to every device in the network. The way the Smurf Attack works is by spoofing the IP address of an ICMP echo request with a victim's IP address and sending it to the broadcast address of an amplifying network. The broadcast address will forward this packet to each device in its network, which will all respond to the spoofed address. If the number of devices in the network is high, this will lead to a high amplification factor; sending a relatively small number of request packets to the broadcast address will then theoretically be able to overwhelm a victim with traffic from the reflected replies.

We implemented the mechanism for a Smurf Attack by modifying our echo request script to account for broadcast addresses. To calculate the broadcast address, since that cannot be determined solely from the IP address, we applied a bitwise-OR with the inverses of the three base subnet masks of the most common network classes (/8, /16, and /24) [5]. We constructed a set of all broadcast addresses (to remove duplicates) before sending echo requests so that we do not send requests to the same network more than once during each measurement.

We also edited our capture analysis script to account for broadcast addresses. This came from the observation that the devices in the network that are sending the responses are not going to be the same as the broadcast address we calculated. Instead, they will be other addresses that map to the same broadcast address. Thus, when measuring the amplification factor for a broadcast address, we had to count any packet sent from an IP that mapped to the current broadcast address we were filtering on, rather than an exact match.

3.4 Ethical Considerations and Limitations

We did not want to overload any individual server with requests, so we made sure we did not send requests too frequently. We ran our scripts sparingly and with at least five minutes of buffer time between each run. Additionally, we did not spoof our IP address, electing instead to receive packets at the same IP that we sent them from. This adds a limitation that our amplification factors are valid only on unspoofed traffic so it does not guarantee that an adversary who spoofs their ICMP traffic gets the same amplification as us.

We also decided to generate our broadcast addresses based on classful addressing which does not reflect ownership over partitioned IPv4 address blocks in the modern Internet. We did this because it was infeasible to ping every possible broadcast address of 500,000 IP addresses.

4 Results

After we ran our measurements, we found that a small proportion of IP broadcast addresses with AF > 1. We created three separate result files filtered by the subnet mask that we applied to the IP addresses we scanned (i.e. by each class of network we sent ICMP packets to). We found that the broadcast addresses from Class A networks were completely unaffected by the Smurf attack, having all generated an AF $<= 1$. This is not surprising since the routers of a Class A network would typically be operated by large corporations and other powerful organizations that have a vested interest in protecting their networks from simple attacks. The routers of Class B and C networks would generally belong to smaller corporations or individual network operators, who have less of an incentive to protect their servers or may not even be aware of such attack patterns. Figures 2 and 3 contain histograms, where the x-axis represents a given AF and the y-axis represents how frequently we saw it. As expected, there are significantly more Class C networks with AF > 1 on the /24 subnet mask (801 total, compared to 4 on /16 and 0 on /8). Most of the networks had low AF's, with a roughly inverse-exponential distribution. Among Class C networks, we also found that 68 of them had a large AF (we follow Moon et al.'s precedent by defininining this as an AF > 10), which is 8.49% of all potentially amplifying networks.

The four Class B networks that generated AF > 1 belonged to Amazon, Universidade Federal do Rio Grande do Sul, SK Broadband (a Korean ISP), and Columbia University. The highest amplification factor in this group is 3, implying that none of them are vulnerable enough to Smurf Attacks to be of use to attackers. However, an

Table 1: Summary of vulnerable Class B and Class C networks with AF > 1

	Median	Mean	Highest	Total
Class B (/16)	2	2.5	3	4
Class C (/24)	3	4.82	62	801

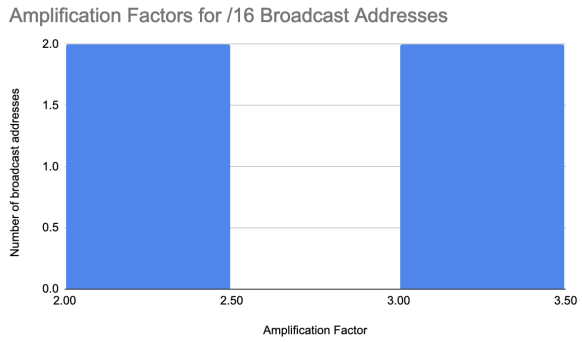

Figure 2: Amplification Factor Histogram of vulnerable Class B networks

AF > 1 does indicate there is some intermediary router that broadcasts IP-directed ICMP packets to more than one network/node. We know this router cannot be an end router because the corresponding Class C networks do not generate an AF > 1.

The highest amplification factor of a Class C network is 62 and corresponds to Comcast (ASN = 7922). Since Comcast is a large ISP, we may assume that there is a router owned by either the ISP or an individual customer in their network that does not disable IP-directed broadcasts. The other Class C networks with high AFs are Conterra (AF = 57, ASN = 32505) and Telstra (AF = 56, ASN = 1221, an Australia-based ISP).

From the geographical distributions, we could reason that all the vulnerable Class B and Class C networks were not anomalously concentrated in any country.

4.1 Response Timestamps

While performing our experiment, we were pinging around 500,000 unique broadcast addresses in less than an hour. Our script to analyze tshark capture files did not take response time into account so we manually looked through the timestamps of the ICMP replies to make sure there was no anomalous delay between the request timestamp and each response timestamp. We looked at a random sample of 10 broadcast addresses from the 801 amplifying Class C networks and found almost no delay between the time the request packet was sent out and the responses. The responses themselves only had a fraction of a millisecond between each response's capture log entry in tshark. For example, in Figure 6, the y-axis represents the amount of milliseconds after the single request packet to a Class C network is sent out.

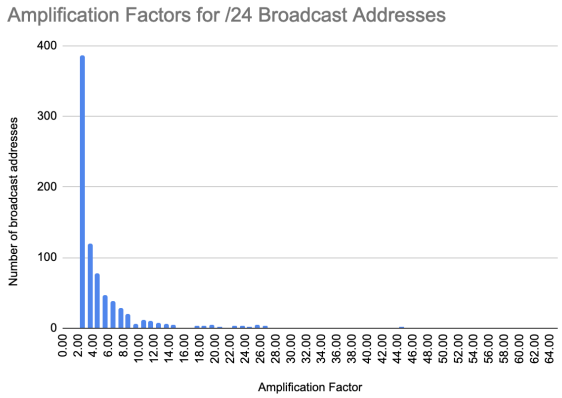

Figure 3: Amplification Factor Histogram on vulnerable Class C networks

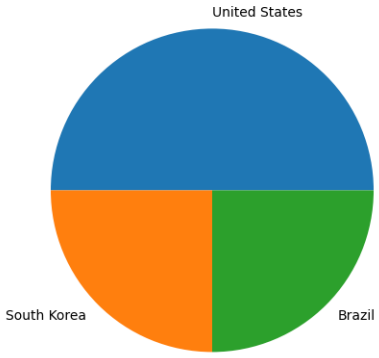

Figure 4: Geographical Distribution of 5 Class B networks that have AF > 1

The almost non-existent delay indicates that the amplified response to a Smurf Attack is ideal for DoS because the responses arrive in almost the same instant, so they will be more effective in overwhelming a specific network or machine.

4.2 Other Attack Surfaces

Another possible attack surface for ICMP-based amplification is known vulnerabilities in software that allow for certain servers to amplify traffic.

One such CVE, published in December 1999, is CVE-2000-0041, which causes Macintosh systems to generate large ICMP datagrams in response to malformed packets [2]. This allows them to be used as an amplifier. We

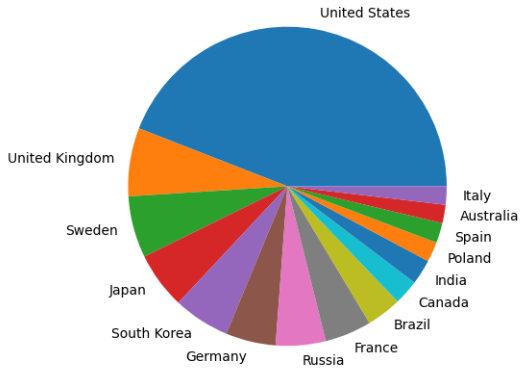

Figure 5: Geographical Distribution of the 801 Class C networks that have AF > 1

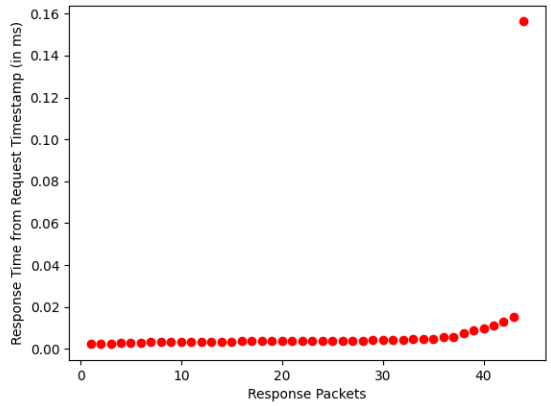

Figure 6: Response Timestamps of a Class C network with AF $= 44$

surveyed whether any public servers used affected cpe configurations (cpe:2.3:o:apple:macos:9.0:*:*:*:*:*:*:*:*) using Censys.io but found none [11].

Another CVE is CVE-1999-1201 which affects Windows 95 and 98 systems [1]. When TCP/IP stacks are bound to the same MAC address, it causes several ping responses upon receiving an ICMP echo packet or TCP SYN. This can also be used to cause TCP Chorusing where multiple TCP ACKs are sent out after receiving a single SYN. A simple solution is keeping only one IP stack bound to a single MAC address on Windows 95 and 98 systems. Using Censys [11], we found 292 remote hosts that run Windows 95 and none that ran the affected cpe configuration for Windows 98. Out of the potentially dangerous 292 remote hosts, 275 are from Amazon-02 (ASN = 275) and 17 from Amazon-AES (ASN = 17). Each of these hosts runs many services for an average of 45.98 TCP-based services on each host. 97% of all these services are HTTP. Since these hosts are from Amazon, it suggests that these 292 hosts are operated by AWS customers who probably bound multiple IP stacks on a single machine, suggesting that these hosts could be used as amplifiers where the amplification factor depends on the number of IP stacks on each MAC address.

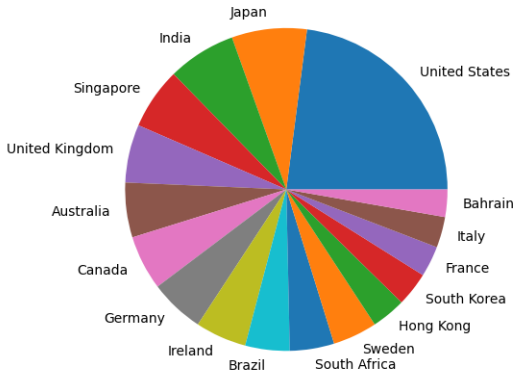

Figure 7: Geographical Distribution of 292 hosts possibly vulnerable to CVE-1999-1201. This data is from Censys.io

There are other possible CVEs rooted in ICMP implementations that cause Denial of Service. A recent example is CVE-2021-40114 [3] where multiple Cisco products have improper memory resource management when processing ICMP traffic. An attacker could send a series of ICMP packets, exhaust resources and crash a vulnerable device. Such a CVE does not fit our requirements since it does not involve ICMP-based amplification.

5 Discussion

Comparing our results to a similar study, Moon et al. [16] measured the AF on several other UDP-based protocols including NTP, SNMP, and SSDP. They were able to find higher median and maximum amplification factors for these protocols than we did for ICMP, and a higher proportion of their vulnerable servers had AF > 10. For NTP, they found several servers (around 700) with AF over 100; for SNMP, they constructed attack patterns that had mean amplification factors of 22.4, 31.8, and 35; and for SSDP, they were able to find "many" servers with AF > 10. This may mean that ICMP-based amplification attacks are not effective as UDP-based ones. However, this does not mean that ICMP-based amplification attacks are impossible. The existence of Class B and Class C networks that have AF > 1 indicate that

there are still flawed network filtering rules in intermediary/end routers. Since these networks are probably partitioned in much smaller IPv4 blocks that are operated by individuals/corporations, it is hard for network operators to coordinate and figure out which intermediary router is facilitating the Smurf attack by forwarding broadcast-directed echo requests.

On the geographical distributions of potential amplifiers, we were not able to find any patterns. Most IPv4 addresses are controlled by routers in the United States so it is normal for it to dominate the distribution of a subset of all IPv4 servers.

6 Conclusion

By pinging the broadcast addresses of Class A, Class B and Class C networks of 500,000 IP addresses, we successfully simulated an over-twenty-year-old attack pattern, the Smurf Attack, without IP spoofing. By measuring the Amplification Factor (ratio of response size and query size) of these addresses, we found 805 vulnerable networks that amplified ICMP traffic on varying levels. However, their amplification factors were not as high as other vulnerable amplifiers to UDP traffic [8, 16]. We also found 292 remote hosts that use outdated OS software susceptible to ICMP/TCP amplification [2] which could be solved by configuring it differently or updating the OS version. Therefore, our results verify there is minimal, but not non-existent, risk for ICMP-based amplification compared to UDP-based amplification. With regards to Smurf Attack mitigation, this risk may be hard to eliminate completely due to the multiple layers of defense required.

References

[1] CVE-1999-1201. Available from MITRE, CVE-ID CVE-1999-1201., 1999.

[2] CVE-2000-0041. Available from MITRE, CVE-ID CVE-2000-0041, 1999.

[3] CVE-2021-40114. Available from MITRE, CVE-ID CVE-2021-401141., 2021.

[4] Icmp type and code ids - ibm documentation. `https://www.ibm.com/docs/en/qsip/7.4?topic=applications-icmp-type-code-ids`, Feb. 2022.

[5] Tcp/ip addressing and subnetting - windows client — microsoft docs. `https://docs.microsoft.com/en-us/troubleshoot/windows-client/networking/tcpip-addressing-and-subnetting`, Feb. 2022.

[6] AKAMAI. Memcached ddos explained. https://www.akamai.com/our-thinking/threat-advisories/memcached-ddos-explained, 2018.

[7] ANTONAKAKIS, M., APRIL, T., BAILEY, M., BERNHARD, M., BURSZTEIN, E., COCHRAN, J., DURUMERIC, Z., HALDERMAN, J. A., INVERNIZZI, L., KALLITSIS, M., ET AL. Understanding the mirai botnet. In *26th USENIX security symposium (USENIX Security 17)* (2017), pp. 1093–1110.

[8] BOCK, K., ALARAJ, A., FAX, Y., HURLEY, K., WUSTROW, E., AND LEVIN, D. Weaponizing middleboxes for {TCP} reflected amplification. In *30th USENIX Security Symposium (USENIX Security 21)* (2021), pp. 3345–3361.

[9] BOUYEDDOU, B., HARROU, F., SUN, Y., AND KADRI, B. Detection of smurf flooding attacks using kullback-leibler-based scheme. In *2018 4th International Conference on Computer and Technology Applications (ICCTA)* (2018), IEEE, pp. 11–15.

[10] CLOUDFLARE. Smurf ddos attack. https://www.cloudflare.com/learning/ddos/smurf-ddos-attack/.

[11] DURUMERIC, Z., ADRIAN, D., MIRIAN, A., BAILEY, M., AND HALDERMAN, J. A. A search engine backed by Internet-wide scanning. In *22nd ACM Conference on Computer and Communications Security* (Oct. 2015).

[12] GUNNAM, G. R., AND KUMAR, S. Do icmp security attacks have same impact on servers? *Journal of Information Security 8*, 3 (2017), 274–283.

[13] ISMAIL, S., HASSEN, H. R., JUST, M., AND ZANTOUT, H. A review of amplification-based distributed denial of service attacks and their mitigation. *Computers & Security 109* (2021), 102380.

[14] KUMAR, S. Smurf-based distributed denial of service (ddos) attack amplification in internet. In *Second International Conference on Internet Monitoring and Protection (ICIMP 2007)* (2007), IEEE, pp. 25–25.

[15] KUMAR, S. Smurf-based distributed denial of service (ddos) attack amplification in internet. In *Second International Conference on Internet Monitoring and Protection (ICIMP 2007)* (2007), pp. 25–25.

[16] MOON, S.-J., YIN, Y., SHARMA, R. A., YUAN, Y., SPRING, J. M., AND SEKAR, V. Accurately measuring global risk of amplification attacks using {AmpMap}. In *30th USENIX Security Symposium (USENIX Security 21)* (2021), pp. 3881–3898.

[17] OTHMAN, R. A. R. Understanding the various types of denial of service attack. *Business Week Online* (2000).

[18] RIVAS, W. R., AND KUMAR, S. Evaluation of centos performance under iot based ddos security attacks. In *2020 3rd International Conference on Data Intelligence and Security (ICDIS)* (2020), IEEE, pp. 64–70.