Messaging, Malware and Mobile Anti-Abuse Working Group

M³AAWG Introduction to Reflective DDoS Attacks

May 2017

The reference URL for this document: www.m3aawg.org/Reflective-DDoS-Introduction

Introduction

Distributed Denial of Service (DDoS) attacks are a crucial concern for many businesses today. Many thousands of individual DDoS attacks take place each day, and though most are relatively small (5-10 gigabits per second), they are still more than sufficient to take unprepared sites offline. Moreover, attackers of even relatively modest means can create attacks in the hundreds-of-gigabits-per-second range. These threaten connectivity over large regions of the internet. It is in everyone's interest to take all possible precautions to thwart these damaging attacks.

There are several types of DDoS attacks. This document tackles just one: the **reflective amplification attack**. This paper is not a best practice document as such; its main purpose is to provide an overview of how this very common form of attack works and what measures can be taken to help eliminate it. It also provides pointers to some of the many technical documents that can provide the detail this overview leaves out.

Reflective Amplification Defined

A reflective amplification attack can be compared to a hijacked conference call. The first person, Alice, purposefully misidentifies herself as her intended victim, Vera. Alice directs a short question to a second person, Bob, that requires a long answer, such as a list of relatives' names and addresses. Bob responds to Vera (who did not ask the question) with a very long answer. Alice's short message has been *amplified* and the reply is being *reflected* to Vera. Repeated many times, with Alice impersonating multiple people (Charlie, Deborah, Edna, and so on), these communications will completely swamp Vera with large amounts of unwanted noise.

In the case of a DDoS attack, the message is *amplified* when compromised devices send a short message to a system that responds with a much bigger payload in the answer. It is *reflected* because the IP address of these requests are forged so that all the responses are sent to a targeted victim rather than the originating device.

A Reflective Amplification DDoS Attack

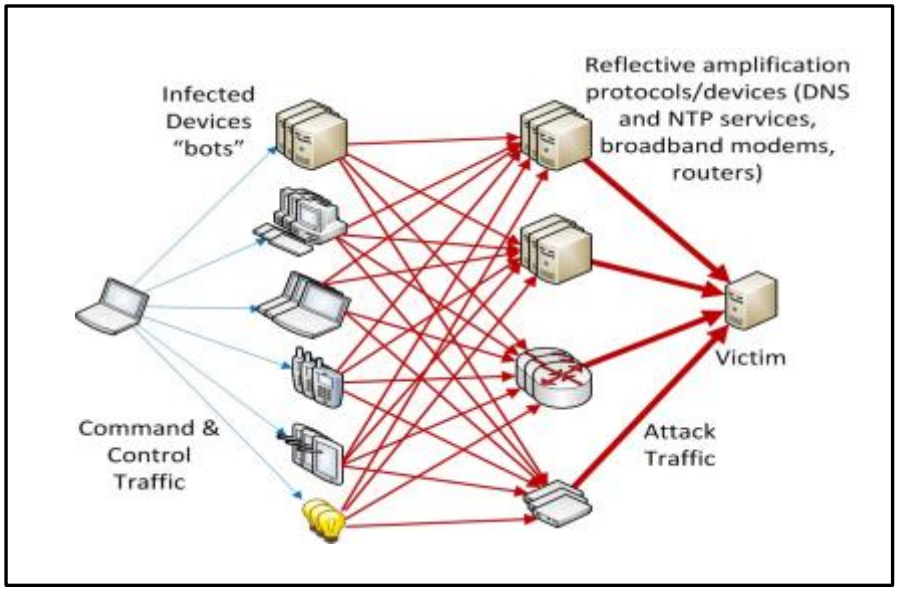

The objective of a reflective amplification attack is to overwhelm a website, server or other network resource with high levels of traffic so that it is unavailable to legitimate users. Below is a list of the common User Datagram Protocols (UDP) and the potential amplification factor for each protocol.

Technical Background: Source Address Spoofing

All Internet Protocol (IP) packets contain two IP addresses. The destination IP address specifies the destination to which the packet will be delivered. The source IP address specifies the destination from which the packet is said to have come. If an answer to the arriving packet is required, that answer will be sent by the destination to the source address.

In Transmission Control Protocol (TCP) used for most of the familiar internet applications such as email and web browsing, connections are opened using a special "three-way handshake" which ensures that each end of the connection can validate the source and destination addresses. However, in the connectionless UDP, the validity of the source address is taken on trust.

UDP is used for services such as the Domain Name System (DNS) where it is appropriate to use a lightweight protocol to get a prompt answer. Send one packet asking for a hostname to be resolved and receive one packet back with the answer and the job is done!

However, if a bad actor were to construct a DNS request using the source IP address of an innocent victim's machine, the answer would not be delivered to the bad actor, but to the victim. If the attacker were to do this at a substantial rate, the victim might not be able to cope with the traffic. They would be suffering a denial-of-service attack using the technique called *reflection*.

If many different DNS servers were used, the attack would be called *distributed* (the "distributed" in DDoS). It can be quite challenging to filter out unwanted traffic and leave only genuine traffic.

Furthermore, a DNS request can be made quite small (the packet header plus a mere handful of bytes), but the answer can be rather large (tens of thousands of bytes in some cases). Thus there can be a very

significant *amplification* factor in play. The attacker need only send out traffic at single-gigabit rates to cause attacks in the tens or perhaps hundreds of gigabits, which will cause outages for most sites.

Besides the DNS protocol, there are other UDP services that can be abused in a similar manner, including NTP (Network Time Protocol) used for obtaining accurate timestamps, SSDP (Simple Service Discovery Protocol) used for provisioning purposes on home networks, and CHARGEN (Character Generator Protocol) that is a 1980s scheme for debugging network links. It is also possible to abuse TCP where no valid connection exists by spoofing packets that appear to be part of the three-way handshake. A US Computer Emergency Readiness Team (US-CERT) Alert lists 15 vulnerable protocols. They cite the work of an independent researcher whose paper provides a thorough analysis of the issue.

Alert (TA14-017A) UDP-Based Amplification Attacks
Source: US-CERT

Protocol	Bandwidth Amplification Factor
DNS	28 to 54
NTP	556.9
SNMPv2	6.3
NetBIOS	3.8
SSDP	30.8
CharGEN	358.8
QOTD	140.3
BitTorrent	3.8
Kad	16.3
Quake Network Protocol	63.9
Steam Protocol	5.5
Multicast DNS (mDNS)	2 to 10
RIPv1	131.24
Portmap (RPCbind)	7 to 28
LDAP	46 to 55

Countermeasures: Source Address Validation (SAV)

Reflective amplification DDoS attacks are by no means the only type of denial-of-service attack. However, they are of special concern because attackers of modest means can create attacks that threaten large parts of the internet. Furthermore, because the original source of the traffic cannot be determined by analyzing packet headers, it is seldom possible to make a prompt identification of the original source of the traffic and make an abuse report.

There is a great deal that can be done to combat reflective amplification of UDP packets. One important strategy is to deal with the reflectors. Operating system vendors and system administrators should turn off

legacy protocols such as CHARGEN that provide little to no value for servers operating on the internet. System administrators should routinely patch servers to prevent attacks from servers running older versions of software such as NTP servers.

However, there are many millions of DNS servers that can be used in amplified reflective attacks and similar numbers of SSDP reflectors. The former may be made secure over time. The latter, which is used by many routers and networks, may require millions of consumers to patch or replace physical devices within their homes, which takes significant time and resources. To more quickly mitigate the SSDP reflector issue, some ISPs have restricted or blocked the SSDP protocol and other commonly abused protocols on their consumer access networks.

Widely available internet services may become reflectors through vulnerabilities, system misconfigurations, improper deployments or for other reasons. When this happens, resolving the issue can take a relatively long time. Reflective DDoS attacks will continue during this remediation phase and must be addressed with more immediate mitigation techniques.

It is sometimes possible to mitigate reflective DDoS attacks by rate-limiting or by blocking the source traffic prior to its being received at the reflector. When the attack traffic has distinctive characteristics, which is often the case, then illicit traffic can just be dropped while allowing legitimate traffic through. Bulk rate-limiting or complete blocking can be used on an emergency basis to stop or reduce the size of the attack while other, less disruptive options are considered. If bulk rate-limiting is necessary, some operators set a limit at twice the normal traffic volumes for affected protocols. This allows some legitimate traffic to keep flowing while reducing the size of the attacks until longer-term remediation is in place.

The most effective – and also the most generalized – way of dealing with reflective DDoS attacks is to make it impossible for packets to be generated with incorrect source IP addresses. This approach is sometimes called "SAV" (Source Address Validation). It is also referred to as "BCP 38," after the IETF's Best Current Practice document BCP 38 (also known as RFC 2827). If SAV were universally applied, it would no longer be possible to mount a reflected attack for any protocol at all because it would no longer be possible to impersonate an IP address. Although this would not get rid of all denial-of-service attacks, it would simplify identifying the sources of attacks and eliminate a complete class of damaging DDoS attacks. Victims would be able to inspect incoming packets and know where they came from, and hence, whom to contact to make the abuse stop.

SAV with spoofed packet dropping is supported in Cable Modem Termination Systems (CMTS) equipment deployed in cable operators access networks globally. This feature became available in the Data Over Cable Service Interface Specification (DOCSIS) release 3.0, first issued in 2006 as a mandatory requirement. The DOCSIS specification requires that SAV be turned on by default for DOCSIS 3.0 and 3.1 compliant CMTS devices.

Source Address Validation Everywhere (SAVE): Problems and Obstacles

There are several different policies and procedures under the SAV/BCP 38 banner, some of which apply to edge systems and some to devices that are more centrally placed within the internet.

For edge systems, the approach can be very simple. Only a small subset of IP addresses, often just one, can be validly used. The ISP offering internet connectivity can discard any packets that use other IP addresses. The equipment used by ISPs to handle customer connections (be they leased line, cable, ADSL or even dial-up) provides the appropriate functionality to enforce source address validation as standard. Once source

address validation has been enabled, the problem is solved from the point of view of the wider internet. However, it is still possible to spoof local addresses. The IETF's Source Address Validation Improvement (SAVI) initiative has documented a number of ways that this further issue can be addressed (see RFC 6959 and RFC 7039).

Multi-homed customers (those who have more than one connection to the internet) can pose a problem if they wish to send packets out of one connection and receive the responses via another. ISPs that are prepared to allow this will need a method of configuring appropriate filters to allow extra ranges of IP addresses to be used. BCP 84 (RFC 3704) documents the issues that arise.

Network providers – companies that transmit internet traffic for others – will see that BCP 38 expects them to install filters that ensure packets with incorrectly-set source addresses are blocked. There is automated support for this with the various flavors of unicast Reverse Path Forwarding (uRPF), details of which can be found in RFC 3704 / BCP 84. In essence, a router checks its forwarding rules to determine whether or not it is plausible for a packet with a particular source address to arrive on a particular interface. If not, the packet is dropped. Although the uRPF approach can be extremely effective, some complex networks with many customers find that making any real difference is often impractical. There is a full discussion of the issues in RFC 3704 / BCP 84.

Hosting companies may allocate IP addresses for their customers to use. Under these circumstances, they have a relatively easy task in knowing which source addresses are valid. Validation may not be much more complex when customers use virtual private servers, even when these come and go depending upon demand.

However, some customers will wish to use their own blocks of IP address space. This will mean that routers have to be configured with a complex set of rules in order to perform source address validation. Some hosting companies use an uRPF approach to dynamically determine whether source addresses are valid. Others prefer to use explicit filters. At hosting companies that have invested in appropriate technology, part of the procedure for bringing a new customer on board is the automated installation of filters. If there is no automation, filters must be installed manually. This can be time consuming and error prone. It is therefore regrettable, but not entirely surprising, that many hosting companies do not currently implement SAV.

Source Address Validation: Economic Disincentives

Economic theory helps explain why source address validation is not yet universally implemented. For *a single individual* to be protected against attacks involving spoofed-source IP addresses, it is necessary for *everyone else* to be disallowing the spoofing. However, as already indicated, hosting companies find it expensive and time-consuming to disallow spoofing and so many do not bother. Equally, medium to large network providers sometimes have a limited understanding of the nature of the traffic flows that cross their links (perhaps from customers of customers) and find that attempts to implement BCP 84 can block legitimate traffic.

So for many entities, SAV is hard to implement, and although it protects other people, there is only a very small benefit for the people who actually do the hard work. In other words, the incentives are not aligned, and for many entities it is apparently economically rational not to bother with SAV.

Nevertheless, some commentators claim that this is not entirely correct and that there already are economic incentives that encourage source address validation. Although the source network of a DDoS attack may not be carrying the huge flows associated with amplification, they are carrying the unamplified traffic and

that has a cost. Additionally, the perpetrators of attacks are criminals – and that criminality may extend to failing to pay their hosting bills.

It is difficult to tell from the outside whether a network does source address validation but some techniques do exist. As these techniques are refined and the results are shared, it is possible that incentives will change. It is possible to imagine a world in which networks that fail to follow best practices and validate source IP addresses will face mounting complaints and ultimately other networks may cease to find it sensible to carry their traffic.

Practical Advice: Deploy uRPF Solutions, Make ACLs Part of Customer On-Boarding

ISPs providing service to residential consumers should find it straightforward to ensure that their customers do not emit packets with source IP addresses they are not entitled to use. The equipment used for DSL, cable and even dial-up connections have standard out-of-the-box options that can be used to make sure that this is so. Thus, the practical advice for ISPs is to identify a suitable configuration and to make regular checks that it is being correctly applied.

Hosting companies can often use one of the various flavors of uRPF to ensure that spoofed traffic cannot be generated by their customers. However, this may not scale to very large numbers of customers. Hosting companies should therefore be making plans to integrate customer on-boarding with automated installation of access control lists (ACLs) into switches and routers. Retrofitting such systems may be far from simple, but the companies who have such systems find them to be extremely effective in preventing abuse and reducing unwanted traffic.

Network providers can sometimes identify locations within their network where ACLs or an uRPF solution can be used without unwanted side-effects. However, this is often problematic. Validation is much easier and simpler to do at the edges of the network. Network providers should consider implementing anti-spoofing techniques for their customers as part of their provisioning process. This is relatively easy to do for single-homed and symmetrically traffic-routed, multi-homed customers. Where the customer network is complex, this type of filtering may be difficult to deploy without blocking legitimate traffic flows. Network providers and their customers should consider adding anti-spoofing language to the contract for the service provided to protect both the provider and the customer.

Conclusion

Reflective amplification DDoS attacks are a serious issue and pose significant risk to service availability on the internet. ISPs, hosting companies and other service providers need to implement, monitor and maintain anti-spoofing techniques in their networks to help minimize and eliminate this class of DDoS attack. For single-homed customers, well established techniques are available to network providers that drop spoofed packets without impact to their customers. Network operators implementing SAV on multi-homed customers must be more cautious. However, for customers with symmetric traffic routing, SAV should have no impact on their legitimate traffic. With the increasing frequency of DDoS attacks on the internet, it is in all operators' interests to implement and police SAV techniques on their networks.

Bibliography and Further Reading

RFC 2827: Ferguson, P. and Senie, D. ,"Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing," BCP 38, RFC 2827, DOI 10.17487/RFC 2827, May 2000, http://www.rfc-editor.org/info/rfc2827

RFC 3704: Baker, F. and Savola, P., "Ingress Filtering for Multihomed Networks," BCP 84, RFC 3704, DOI 10.17487/RFC 3704, March 2004, http://www.rfc-editor.org/info/rfc3704

RFC 6959: McPherson, D., Baker, F., and Halpern, J., "Source Address Validation Improvement (SAVI) Threat Scope," RFC 6959, DOI 10.17487/RFC 6959, May 2013, http://www.rfc-editor.org/info/rfc6959

RFC 7039: Wu, J., Bi, J., Bagnulo, M., Baker, F., and Vogt, C., Ed., "Source Address Validation Improvement (SAVI) Framework," RFC 7039, DOI 10.17487/RFC 7039, October 2013, http://www.rfc-editor.org/info/rfc7039

"Data-Over-Cable Service Interface Specifications DOCSIS 3.0 Security Specification," CM-SP-SEC3.0-I16-160602, 02 June 2016, https://community.cablelabs.com/wiki/plugins/servlet/cablelabs/alfresco/download?id=83328160-1b1a-48ab-8333-8a96544c0db2;1.0

Rossow, Christian. "Amplification Hell: Revisiting Network Protocols for DDoS Abuse." In *Proceedings of the 2014 Network and Distributed System Security (NDSS) Symposium*. https://www.internetsociety.org/sites/default/files/01_5.pdf

United States Computer Emergency Readiness Team (US-CERT). "Alert (TA14-017A) UDP-Based Amplification Attacks." January 17, 2014, rev. Nov 4, 2016, https://www.us-cert.gov/ncas/alerts/TA14-017A

Glossary

Access Control List (ACL)	For network communications, access control lists allow or deny network packets. Matching criteria can involve IP addresses, ports, packet length, header fields and almost any bit contained in the network packet, depending on the capability of the filtering device. Typical network filtering devices include routers, switches and firewalls.
CHARGEN	A protocol defined in IETF RFC 864. Originally developed as a debugging and measurement tool, it is rarely used today and is susceptible to misuse in DDoS attacks.

CERT	Computer Emergency Response Team.
Denial of Service (DoS) Attack	Malicious traffic that attempts to deny access to network, server, or application resources.
Distributed Denial of Service (DDoS) Attack	A DoS attack in which the attack traffic comes from (is distributed across) multiple sources, which might be anything from computers to smartphones to IoT-connected devices.
Domain Name System (DNS)	A critical internet service that translates alphanumeric names to IP addresses.
Internet Service Provider (ISP)	A company or organization that provides internet access to its subscribers.
Internet Protocol (IP)	The main protocol used to deliver packets across the internet.
Internet of Things (IoT)	A term used to describe adding network connectivity to a variety of physical objects for local communication or communication across the internet. Examples of devices include light bulbs, refrigerators, washing machines, home fitness equipment, vehicles, traffic signals, soil moisture sensors and much more. Major categories of IoT devices include consumer, smart cities, industrial, healthcare, government, financial and more.
Network Time Protocol (NTP)	A protocol used to synchronize clocks across the internet and other packet-switched networks.
Reflective Amplification Distributed Denial of Service (DDoS) Attack	A DDoS attack in which the IP packets' source address is changed from the actual sender to the victim's IP address. The IP packets are then sent to a service on the internet that amplifies their effect. The original IP packets are "reflected" off the legitimate service and the response goes to the DDoS attack victim, flooding the victim with IP packets they did not request. Common vulnerable protocols include DNS, NTP, SSDP, SNMP and at least 11 others. Some of the largest DDoS attacks seen on the internet have used this technique. In literature, this type of attack may also be a called a reflective DDoS attack.
Simple Service Discovery Protocol (SSDP)	The **Simple Service Discovery Protocol** (**SSDP**) is a network protocol used by uPnP (Universal Plug and Play) for the advertisement and discovery of network services and device information. It was intended for home and small office networks, not for use on the general internet.
Source Address Validation (SAV)	Validating that the source address on an IP packet is actually from the device that was legitimately assigned that IP address on the internet. IETF RFC 2827 / BCP 38 describes a common method of SAV using ingress filtering. IETF RFC 3704 / BCP 84 discusses ingress filtering for multi-homed networks. IETF RFC 7039 discusses improvements to

	complement ingress filtering to prevent nodes attached to the same IP link from spoofing each other's IP addresses. DOCSIS 3.0 and 3.1 Security Specification documents describe the mandatory source address verification technique used in cable modem termination system equipment.
Source Address Validation Everywhere (SAVE)	A movement to require internet ecosystem players (e.g., ISPs, hosting companies, enterprises) to implement SAV across the entire internet.
Source Address Validation Improvement (SAVI)	A framework outlined in IETF RFC 7039 that states: "Source Address Validation Improvement (SAVI) methods were developed to prevent nodes attached to the same IP link from spoofing each other's IP addresses, so as to complement ingress filtering with finer-grained, standardized IP source address validation."
Transmission Control Protocol (TCP)	A network protocol that runs on top of the IP protocol to provide a reliable method of delivering ordered data packets.
User Datagram Protocol (UDP)	A network protocol that runs on top of the IP protocol to provide packet delivery for loss-tolerant applications. An example application is a broadcast television stream.
Unicast Reverse Path Forwarding (uRPF)	A security technique used to check that a route exists in the route table for the source IP address of a received packet. Several variations of uRPF exist, including strict, feasible and loose. These techniques are described in IETF RFC 3704 / BCP 84.

As with all documents that we publish, please check the M³AAWG website (www.m3aawg.org) for updates to this paper.