

## Distributed Denial of Service (DDoS)

# MS-ISAC Guide to DDoS Attacks

May 2023



# Contents

This Multi-State Information Sharing and Analysis Center (MS-ISAC) document serves as a guide to aid partners in defending against Distributed Denial of Service (DDoS) attacks. This guide is not exhaustive, but it provides direct examples of attack types that U.S. State, Local, Tribal, and Territorial (SLTT) partners reported to the MS-ISAC.

---

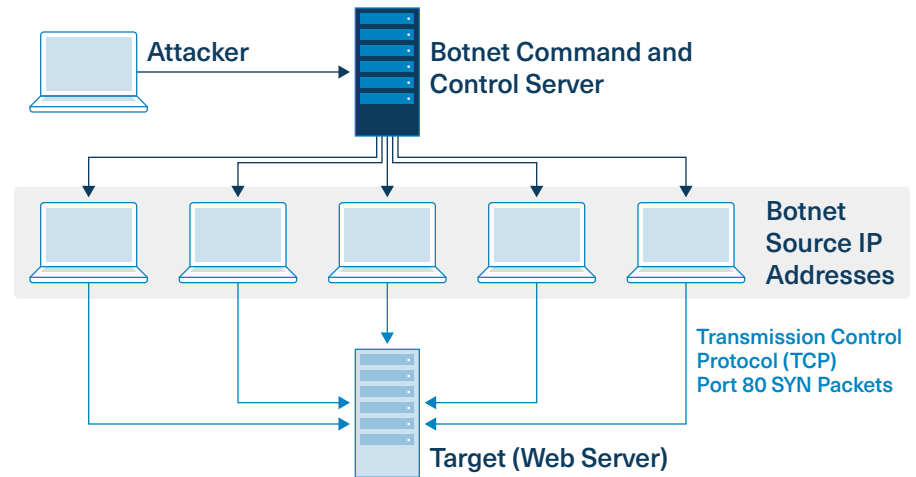
<b>Introduction</b>	<b>1</b>
<hr/>	
<b>Standard DDoS Attack Types</b>	<b>3</b>
SYN Flood	3
UDP Flood	5
SMBLoris	6
ICMP Flood	7
HTTP GET Flood	9
<hr/>	
<b>Reflection DDoS Attack Types</b>	<b>10</b>
NTP Reflection Attack with Amplification	10
DNS Reflection Attack with Amplification	11
CLDAP Reflection Attack with Amplification	12
WordPress Pingback Reflection Attack with Amplification	13
SSDP Reflection Attack with Amplification	14
Microsoft SQL Reflection Attack with Amplification	15
Memcached DDoS Attacks (Amplification)	16
<hr/>	
<b>General Recommendations and Mitigation Strategies</b>	<b>17</b>

# Introduction

A Denial of Service (DoS) attack is an attempt to overwhelm and render a system unavailable to intended user(s), such as preventing their access to a website. A successful DoS attack consumes all available network, application, or system resources, usually resulting in a network slowdown, application crash, or server crash. When multiple sources coordinate in a DoS attack, it is known as a Distributed Denial of Service (DDoS) attack.

MS-ISAC regularly observes two main methods of DDoS attacks: "Standard" and "Reflection."

A standard DDoS attack occurs when Cyber Threat Actors (CTAs) direct substantial network traffic to a target server or network. One of the ways a CTA accomplishes this is by using a botnet to send the network traffic. A botnet is a large number of previously compromised devices (also known as "bots" or "zombies") that can be controlled over the internet from a single location and directed to carry out desired actions. When a CTA uses a botnet to perform a DDoS attack, they send instructions to zombie machines connected to that botnet, thereby magnifying the scale of their attack. By leveraging a botnet, attackers enable a DDoS attack to originate from multiple networks and even multiple countries.




**Figure 1.** Standard DDoS SYN Flood Example

A reflection DDoS attack occurs when attackers spoof their IP address to pose as the intended victim and then send requests to public-facing servers. The responses to these requests are sent to the intended victim from legitimate servers.

In addition to these methods, CTAs use a technique called "amplification" to increase the effectiveness of their attacks. Usually used in conjunction with reflection attacks, amplification occurs when CTAs request large amounts of data from third-party systems to ensure the response sent to the victim is larger than the request sent from the attacker. As Figure 2 illustrates, this might occur when the attacker spoofs its IP address, pretending to be the victim, and requests all known data from a public server. This results in the attacker sending a small request, but the public server responds to the victim with a large amount of data.

CTAs use various techniques to generate the necessary traffic for an effective DDoS attack. A lone CTA with a botnet at their disposal may use that botnet to orchestrate the attacks. However, botnets are also available for hire, with operators charging minimal fees for short-duration attacks.

Rather than trying to gain access to a botnet, a group of CTAs working together may use free tools for a DDoS attack. Attacks like these are usually less successful, as it is difficult to coordinate enough attackers for the effect to be noticeable. Most of these open-source tools were originally designed as stress testers, but lower-skilled CTAs have used them to conduct DDoS attacks. Popular examples of these tools include the Low Orbit Ion Cannon (LOIC) and the High Orbit Ion Cannon (HOIC). These tools can be downloaded, installed, and used by anyone who wishes to be a part of an ongoing DDoS attack. Other tools used to perform DDoS activities include Metasploit, Pyloris, Slowloris, HULK, DDOSIM, Torshammer, and GoldenEye.



**Figure 2.** DNS Reflection DDoS with Amplification Example

# Standard DDoS Attack Types

## SYN Flood

A SYN Flood is one of the most common forms of DDoS attacks observed by the MS-ISAC. It occurs when an attacker sends a succession of TCP Synchronize (SYN) requests to the target in an attempt to consume enough resources to make the server unavailable for legitimate users. This works because a SYN request opens network communication between a prospective client and the target server. When the server receives a SYN request, it responds by acknowledging the request and holds the communication open while it waits for the client to acknowledge the open connection. However, in a successful SYN Flood, the client acknowledgment never arrives, thus consuming the server's resources until the connection times out. A large number of incoming SYN requests to the target server exhausts all available server resources and results in a successful DDoS attack.

## SYN Flood Variations

**Slowloris Attacks** Slowloris is a DoS tool that CTAs can easily access, but it also refers to a specific type of DoS attack. Slowloris attacks attempt to establish multiple TCP connections on a target web server and hold them open for as long as possible by sending partial requests. As such, they are very similar to a SYN Flood.

**ESSYN/XSYN Flood** An ESSYN Flood, also known as an XSYN Flood, is an attack designed to target entities using stateful firewalls. The attack works when a large number of unique source IP addresses all attempt to open connections with the target destination IP. Each new connection from a unique source IP creates a new entry in the firewall state table. This attack aims to create more unique connections than for which there is space in the firewall's state table. Once the table is full, the firewall will no longer accept any additional inbound connections, denying service to legitimate users attempting to access the destination IP.

**PSH Flood** A Push (PSH) Flood involves sending a large number of TCP packets with the PSH bit enabled. The purpose of a PSH packet is to bypass packet buffering, which allows for the efficient transfer of data by ensuring packets are filled to the maximum segment size when multiple packets are sent over a TCP connection. If the PSH bit is enabled, it indicates the packet should immediately be sent to the application. In normal circumstances, this does not present an issue. However, when a significant number of PSH packets are sent to a target server, there is potential to overload its resources, creating a DoS situation.

## Recommendations

To identify a SYN Flood, investigate network logs and locate the TCP SYN flag. Tcpdump or Wireshark may work for this purpose.

- TCP SYN packets are normal and not necessarily indicative of malicious activity. Instead, look for a large number of SYN packets from multiple sources over a short duration.

If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

To help minimize the impact of successful SYN Flood attacks, define strict “TCP keepalive” and “maximum connection” rules on all perimeter devices, such as firewalls and proxy servers.

- On some firewall appliances, you can enable “SYN cookies” to help mitigate the effects of a SYN Flood. Enabling SYN cookies forces the firewall to validate the TCP connection between the client and server before traffic is passed to the server. When attackers never send a final acknowledgment of the open connection, the firewall drops the connection.

## UDP Flood

A UDP Flood is very similar to a SYN Flood in that an attacker uses a botnet to send a significant amount of traffic to the target server. The difference is that this attack is much faster. Rather than attempting to exhaust server resources, it seeks to consume all of the available bandwidth on the server's network link, thereby denying access to legitimate users. The attack works because a server receiving a UDP packet on a network port, such as 50555/UDP, checks for an application listening on that port. If nothing is listening on the port, it replies to the sender of the UDP packet with an Internet Control Message Protocol (ICMP) Destination Unreachable packet. During an attack, a large number of UDP packets arrive, each with various destination ports. This forces the server to process and, in most cases, respond to each one. This attack can quickly consume all available bandwidth.

### Recommendations

To identify a UDP Flood, investigate network logs and look for a large number of inbound UDP packets over irregular network ports coming from a large number of source IP addresses.

- Many legitimate services use UDP for their network traffic. Common UDP ports are 53 (DNS), 88 (Kerberos), 137/138/445 (Windows), and 161 (SNMP). When investigating a DDoS attack, look for UDP traffic with high-numbered network ports (1024+).

If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

- To minimize the effect of UDP Flood attacks, define strict rules on your perimeter network devices, like firewalls, to allow only inbound traffic on required ports.

A Server Message Block (SMB) SMBLoris attack is an application-level DDoS attack that occurs when a CTA opens multiple SMB connections to a device, maliciously consuming memory with minimal attack cost. SMB is a remote access protocol that provides shared access to files, printers, and various communications between devices over port 445. All versions of SMB are vulnerable to SMBLoris because the vulnerability lies in how SMB packets are processed and the memory allocated. Windows and Samba software devices are both susceptible to the attack.

A connection made with a single IPv4 or IPv6 address impacts up to 8 GB of memory if a CTA sends the attack over both IPv4 and IPv6, enabling one computer to cause 16 GB of memory to be consumed while utilizing only 512 MB of its own memory. Eventually, a targeted Windows computer cannot allocate any more memory, so it becomes unresponsive and needs to be manually rebooted. If this attack occurs against a Linux device with Samba, the device is forced into its configured Out of Memory (OOM) behavior.

As of April 2023, the MS-ISAC has not received recent reporting on this technique but is including it in this guide in the event this type of activity resumes.

### Recommendations

- To block a remote SMBLoris attack from occurring, configure the border firewall to block all ingress traffic over ports 445 and 139.
- To block an internal SMBLoris attack from occurring, set an artificial rate limit for the number of connections local devices can have open.



## ICMP Flood

An ICMP Flood occurs when an attacker uses a botnet to send a large number of ICMP packets to a target server to consume all available bandwidth and deny legitimate users access.

This attack works when a large number of sources can send enough ICMP traffic to consume all available bandwidth of the target's network.

An example of this could be the "ping" command. This command is primarily used to test network connectivity between two points on a network. However, it is possible to supply this command with different variables to expand the ping's size and frequency. By using these variables correctly and with enough source machines to initiate the traffic, it is possible to consume all of the available bandwidth.

### ICMP Flood Variant Using Reflection

**Smurf Attack** A Smurf attack is an alternate method of carrying out an ICMP Flood attack. In a Smurf attack, the attacker spoofs the target's IP address as their own and then sends ICMP ping requests to the broadcast IP address of a public network on the internet. The broadcast IP address of a network will send any traffic it receives to all other IP addresses within its network. Therefore, when the broadcast IP address receives the ICMP ping request, it is forwarded to all live computers on its network. Each of those computers thinks that these ping requests are coming from the target IP address and therefore sends their responses to the target rather than back to the

attacker. This sends many unsolicited ICMP ping replies to the target of the DDoS and consumes available bandwidth.

### Recommendations

To identify an ICMP Flood, investigate network logs and look for a significant amount of inbound ICMP traffic from a large number of sources.

- Depending on what tool you are using to investigate your logs, you can identify ICMP packets by the protocol displayed in the graphical user interface, such as with WireShark. When analyzing ICMP traffic, you will notice that no port information is available, as ICMP does not use network ports like TCP or UDP.
- If you are using a tool that displays the network protocols as numbered values, ICMP is protocol 1.

ICMP type and code fields also identify what ICMP traffic is being sent or received. For a complete list of these types and codes, please see <http://www.nthelp.com/icmp.html>.

- If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

To mitigate some of the damage of ICMP Flood attacks, block ICMP traffic at perimeter network devices such as routers.

Additionally, set a packet-per-second threshold for ICMP requests on perimeter routers. If the amount of inbound ICMP traffic exceeds this threshold, the excess traffic is ignored until the next second. Packet-per-second thresholds effectively keep your network from being overrun with ICMP traffic.

- Note: The above step does not stop a determined ICMP Flood. If there is enough inbound traffic to exhaust the bandwidth between the upstream network provider and the perimeter device-filtering ICMP, legitimate traffic may be dropped or delayed to the point of a DoS. If this is the case, it is necessary to contact the upstream network service provider to have ICMP activity dropped at their level before it reaches your network link.

## HTTP GET Flood

An HTTP GET Flood occurs when an attacker, or attackers, generates a significant number of continuous HTTP GET requests for a target website in an attempt to consume enough resources to make the server unavailable for legitimate users. In this case, the attacking IP addresses never wait for a response from the target server despite the server attempting to respond to all incoming requests. This results in connections being left open on the web server. A large enough number of incoming HTTP GET requests to the target web server eventually exhausts all available server resources and results in a successful DDoS attack.

### HTTP GET Flood Variation

**HTTP POST Flood** Another HTTP Flood incorporates the use of the HTTP POST request instead of GET. This attack works because it forces the web server to allocate more resources in response to each inbound request. A large number of these requests could tie up enough server resources to deny legitimate users access to the web server.

### Recommendations

To identify an HTTP GET Flood, investigate network logs and look for a large number of inbound traffic from a significant number of source IP addresses with a destination port of 80 and a protocol of TCP. The packet data should also begin with "GET." We recommend using either tcpdump or Wireshark.

HTTP GET requests are normal and are not inherently indicative of malicious activity. Look for a large number of identical GET requests coming from a large number of sources over a short period. The same source IP addresses should resend the same GET requests rapidly.

- If you identify an attack, leverage a DDoS mitigation service provider for the best results in mitigating this activity.
- It is difficult to set up proactive security measures to block this attack, as legitimate traffic is used to carry it out. Often, rate-based protections are not sufficient to block this attack, and the source IP addresses of the attack are part of a large botnet. Therefore, blocking every source IP is inefficient and may include legitimate users.
- One solution that may help mitigate this type of attack is to use a Web Application Firewall (WAF). HTTP Floods often exhibit trends that a correctly configured WAF can filter and block without blocking legitimate access to the web server.

# Reflection DDoS Attack Types

## NTP Reflection Attack with Amplification

A Network Time Protocol (NTP) reflection attack occurs when the attacker uses traffic from a legitimate NTP server to overwhelm a target's resources. NTP synchronizes clocks on networked machines and runs over port 123/UDP. An obscure command, monlist, allows a requesting computer to receive information regarding the last 600 connections to the NTP server. An attacker can spoof the target's IP address and send a monlist command to request that the NTP server send a large amount of information to the target. These responses typically have a fixed packet size that can be identified across a large number of replies. Since the response from the NTP server is larger than the request sent from the attacker, the effect of the attack is amplified. When an attacker spoofs the target's IP address and then sends the monlist command to a large number of internet-facing NTP servers, the amplified responses are sent back to the target. This eventually consumes all available bandwidth.

### Recommendations

To identify a NTP Reflection Attack with Amplification, investigate your network logs and look for inbound traffic with a source port of 123/UDP and a specific packet size.

- Once identified, try to leverage your upstream network service provider and provide them with the attacking IP addresses and the packet sizes used in the attack. Upstream providers have the ability to place a filter at their level to force inbound NTP traffic, using the specific packet size that you are experiencing, to drop.

Along with remediating inbound attacks, take the following preventative measures to ensure that your NTP servers are not used to attack others.

- If you are unsure whether or not your NTP server is vulnerable to attack, follow the instructions available at OpenNTP: <http://openntpproject.org>.
- Upgrade NTP servers to version 2.4.7 or later, which removes the monlist command entirely, or implement a version of NTP that does not utilize the monlist command, such as OpenNTPD.
- If you are unable to upgrade your server, disable the monlist query feature by adding "disable monitor" to your ntp.conf file and restarting the NTP process.
- Implement firewall rules that restrict unauthorized traffic to the NTP server.

## DNS Reflection Attack with Amplification

A Domain Name System (DNS) Reflection attack occurs when the attacker manipulates the DNS system to send an overwhelming amount of traffic to the target. DNS servers resolve IP addresses to domain names, allowing the average internet user to type an easily remembered domain name into their web browser rather than remembering the IP addresses of websites. A DNS Reflection attack occurs when an attacker spoofs the victim's IP address and sends DNS name lookup requests to public DNS servers. The DNS server then sends the response to the target server, and the response size depends on the options specified by the attacker in their name lookup request. To get the maximum amplification, the attacker can use the word "ANY" in their request, which returns all known information about a DNS zone to a single request. When an attacker spoofs a target's IP address and sends DNS lookup requests to a large number of public DNS servers, the amplified responses are sent back to the target and eventually consume all available bandwidth.

### Recommendations

To identify if a DNS Reflection Attack with Amplification is occurring, investigate network logs and look for inbound DNS query responses with no matching DNS query requests.

- DNS queries are normal and are themselves not indicative of an attack.

If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

Along with remediating inbound attacks, disable DNS recursion, if possible, by following the guidelines provided by your DNS server vendor (BIND, Microsoft, etc.). Doing so ensures that your DNS servers are not used to attack others.

- Instructions for disabling recursion can also be found at Team Cymru: <http://www.team-cymru.org/Services/Resolvers/instructions.html>.

To discover if any of your public DNS servers may be used to attack others, use the free test at [openresolverproject.org](http://openresolverproject.org).

## CLDAP Reflection Attack with Amplification

A Connection-less Lightweight Directory Access Protocol (CLDAP) Reflection Attack with Amplification occurs when an attacker sends a CLDAP request to an LDAP server using a spoofed sender IP address. CLDAP is used to connect, search, and modify shared internet directories. It runs over port 389/UDP.

A CLDAP Reflection attack occurs when a CTA spoofs the victim's IP address and sends a CLDAP query to multiple LDAP servers. The LDAP servers then send the requested data to the spoofed IP address. This unsolicited response results in a DDoS attack, as the victim's machine cannot process an overabundance of LDAP/CLDAP data simultaneously.

The amplification is due to the number of times a packet is enlarged while processed by the LDAP server. LDAP UDP protocol responses are much larger than the initial request with an amplification factor of 52, and they can peak at up to a factor of 70.

### TCP LDAP Reflection Attack with Amplification Variant

The LDAP Reflection Attack with Amplification variant can be used over port 389/TCP. This attack has an amplification factor of 46 and can peak at up to a factor of 55.

## Recommendations

To identify a CLDAP Reflection Attack with Amplification, investigate your network logs and look for inbound traffic with a source port of 389/UDP.

- Once identified, try to leverage your upstream network service provider and provide them with the attacking IP addresses and the packet sizes used in the attack. Upstream providers can place a filter at their level.
- Create a DDoS protection plan.

Along with remediating inbound attacks, take the following preventative measures to ensure that your servers are not used to attack others.

- Implement ingress firewall rules that restrict unauthorized use of the LDAP server.
- Audit policies to provide reporting of network services that are potentially exploitable as reflection attacks.

## WordPress Pingback Reflection Attack with Amplification

WordPress is a popular Content Management System (CMS) used to develop and maintain websites and blogs. A function of WordPress sites is called the pingback feature, which notifies other WordPress websites that you have added a link to their website on your site. Sites using WordPress automate this process and maintain automated lists containing sites that link to them. These “pingbacks” are sent as Hypertext Transfer Protocol (HTTP) POST requests to the /xmlrpc.php page, which WordPress uses to carry out the pingback process. By default, this feature downloads the entire web page that contains the link that triggered the pingback process. An attacker can locate any number of WordPress websites and then send pingback requests to each of them with the URL of the target website, resulting in each of those WordPress websites sending requests to the target server requesting the download of the web page. A large number of requests to download the web page can eventually overload the target web server.

### Recommendations

To identify a WordPress Pingback Reflection attack with Amplification, investigate your network logs and look for a large number of inbound TCP traffic over port 80 from a large number of sources. The traffic appears as HTTP GET requests for random values such as “?5454545=6767676”. This request bypasses the cache and forces a full-page reload for every packet.

- If you identify an attack, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

At the time of this writing, there is no way to prevent this inbound traffic, as it functions as normal web traffic. However, there is a way to ensure that your WordPress websites are not used to attack others. To do this, WordPress has a tool available for download that disables the pingback feature of XMLRPC. Download the tool at the following link: <https://wordpress.org/plugins/disable-xml-rpc-pingback>.

Alternatively, you can create a plugin for the website that add a filter to manually disable the pingback function of XMLRPC. An example of this plugin can be found at: <https://isc.sans.edu/forums/diary/Wordpress+Pingback+DDoS+Attacks/17801>.

## SSDP Reflection Attack with Amplification

The Simple Service Discovery Protocol (SSDP) is commonly used to discover Universal Plug and Play (UPnP) devices. UPnP is a series of networking protocols that allows networking devices to discover and connect with one another without user intervention. Using SSDP, Simple Object Access Protocol (SOAP) delivers control messages to UPnP devices. An SSDP reflection attack occurs when an attacker spoofs the victim's IP address and sends crafted SOAP requests to open UPnP devices on the internet. These devices then send their responses to that victim's IP address. Depending on how the attacker crafted the request, the response could be amplified by a factor of 30 from a single request.

When an attacker spoofs a victim's IP address and sends crafted SOAP requests over SSDP to a large number of public UPnP devices, the amplified responses are sent back to the victim, eventually consuming all available bandwidth.

### Recommendations

To identify if an SSDP Reflection Attack with Amplification is occurring, investigate network logs and look for inbound source port 1900/UDP (SSDP) traffic from a large number of source IP addresses.

- Once an attack is identified, try to leverage your upstream network service provider in order for them to mitigate the activity before it reaches your network.

Along with remediating inbound attacks, take the following preventative measures to ensure that your UPnP devices are not used to attack others.

- If you are unsure if any devices on your network could be employed in an attack, check by following the instructions available at OpenSSDP: <http://openssdpproject.org>.
- Block outbound port 1900/UDP traffic at your border routers and restrict UPnP to the internal network, if required.



## Microsoft SQL Reflection Attack with Amplification

Microsoft (MS) Structured Query Language (SQL) is a popular application for managing relational databases. Database servers using MS SQL are sometimes left on external IP addresses so that they can be accessed remotely over the internet. An MS SQL Reflection Attack occurs when an attacker spoofs the target's IP address and sends crafted requests to public-facing MS SQL servers using the MS SQL Server Resolution Protocol (MC-SQLR), which listens on port 1434/UDP. The response from the database server contains information about the database instances running on the server and how to connect to each one. Depending on the configuration of the database server and the number of database instances on the server, the response to the client request could be amplified by a factor of 25 for a single request.

Attackers can send scripted MC-SQLR requests, spoofing the target's IP address, to a large number of public-facing MS SQL servers. The amplified responses are sent back to the target, possibly consuming all of the target's available bandwidth.

### Recommendations

To identify if an MS SQL Reflection Attack with Amplification is occurring, investigate network logs and look for inbound source port 1434/UDP (MC-SQLR) traffic from a large number of source IP addresses. In some instances, it may be possible to identify a particular payload signature.

- If you identify an attack, try to leverage your upstream provider in order for them to mitigate the activity before it reaches your network.
- If possible, block inbound connections to port 1434/UDP or filter connections to allow only trusted IP addresses.

Along with mitigating inbound attacks, take the following steps to prevent your MS SQL server from being used as a reflector in attacks against others:

- Use ingress and egress filters on firewalls to block SQL server ports. Port 1434/UDP should be open only if there is an identified need for the service. If the port is open, it is recommended that traffic be filtered to allow only trusted IP addresses.
- SQL servers that have only one database instance running do not need to run MS-SQLR. If you are running only one database instance, disable MS-SQLR.

As of Microsoft SQL Server 2008, the feature is disabled by default. However, earlier versions require administrators to disable this service manually. If you are running an older version of the software and there is no need for MS-SQLR, disable it. If it is determined that MS-SQLR is needed, consider adding an additional layer of security, such as requiring authentication via SSH or VPN, in front of the service.

## Memcached DDoS Attacks (Amplification)

A Memcached DDoS Attack attempts to overload a targeted victim with traffic using the UDP protocol. The attacker spoofs requests to a vulnerable UDP memcached server, which then floods a targeted victim with internet traffic. These requests potentially overwhelm the victim's resources, thus resulting in a denial of service on the server itself. The attack works by sending spoofed requests to the vulnerable server, which then responds with a larger amount of data than the initial request. This increases the overall traffic.

A Memcached DDoS Attack occurs in four steps:

- 1** An attacker implants a large amount of data on an exposed memcached server.
- 2** Next, the attacker spoofs an HTTP GET request with the IP address of the targeted victim.
- 3** The vulnerable memcached server that receives the request, which is trying to be helpful by responding, sends a large response to the target that the target won't be able to handle.
- 4** The targeted server or its surrounding infrastructure is unable to process the large amount of data sent from the memcached server, resulting in an overload and denial of service to legitimate requests.

## Recommendations

- Disable UDP for memcached servers.
- Restrict access to memcached servers from external sources by utilizing a firewall configuration that limits access.

# General Recommendations and Mitigation Strategies

The following generic recommendations for DDoS mitigation can reduce the impact of attempted DDoS attacks and enable a faster response when successful DDoS attacks occur.

- Establish and maintain effective partnerships with your upstream network service provider and know what assistance they can provide you in the event of a DDoS attack. In the case of a DDoS attack, the faster a provider can implement traffic blocks and mitigation strategies at their level, the sooner your services will become available for legitimate users.
- Consider also establishing relationships with companies that offer DDoS mitigation services.
- If you are experiencing a DDoS attack, provide the attacking IP addresses to your upstream network service provider so they can implement restrictions at their level. Keep in mind that Reflection DDoS attacks typically originate from legitimate public servers. It is important to ascertain to whom an IP belongs when examining network logs during an attack. Use tools such as the American Registry for Internet Numbers (ARIN) (<https://www.arin.net>) to look up the source IPs involved in the attack. Otherwise, you may block traffic from legitimate networks or servers.
- Enable firewall logging of accepted and denied traffic to determine where the DDoS may be originating.
- Define strict “TCP keepalive” and “maximum connection” on all perimeter devices, such as firewalls and proxy servers. This recommendation assists with keeping SYN Flood attacks from being successful.
- Consider having the organization’s upstream network service provider implement port and packet size filtering.
- Establish and regularly validate public-facing websites’ baseline traffic patterns for volume and type.
- Apply all vendor patches after appropriate testing.
- Configure firewalls to block, at a minimum, inbound traffic sourced from IP addresses that are reserved (0/8), loopback (127/8), private (RFC 1918 blocks 10/8, 172.16/12, and 192.168/16), unassigned DHCP clients (169.254.0.0/16), multicast (224.0.0.0/4) and otherwise listed in RFC 5735. This configuration should also be requested at the ISP level.
- Tune public-facing server processes to allow the minimum amount of processes or connections necessary to conduct business effectively.
- Configure firewalls and intrusion detection/prevention devices to alarm on traffic anomalies.

- Configure firewalls only to accept traffic detailed in your organization's security policy as required for business purposes.
- Consider setting up Out-of-Band access, internet and telephony, to an incident management room to ensure connection in the event of a DDoS attack that disrupts normal connectivity.
- Consider utilizing a free service such as the Athenian Project. The Athenian Project is a Cloudflare-run program offered free to all U.S state, county, and municipal government websites. They aim to provide protection and reliability for websites from cyber attacks. For more information, please see: <https://www.cloudflare.com/athenian>.
- Ensure all software is up to date, as vulnerabilities that could allow your servers to be used for attacks against other victims are often patched by software updates.

The Center for Internet Security, Inc. (CIS®) makes the connected world a safer place for people, businesses, and governments through our core competencies of collaboration and innovation. We are a community-driven nonprofit, responsible for the CIS Critical Security Controls® and CIS Benchmarks™, globally recognized best practices for securing IT systems and data. We lead a global community of IT professionals to continuously evolve these standards and provide products and services to proactively safeguard against emerging threats. Our CIS Hardened Images® provide secure, on-demand, scalable computing environments in the cloud.





CIS is home to the Multi-State Information Sharing and Analysis Center® (MS-ISAC®), the trusted resource for cyber threat prevention, protection, response, and recovery for U.S. State, Local, Tribal, and Territorial government entities, and the Elections Infrastructure Information Sharing and Analysis Center® (EI-ISAC®), which supports the rapidly changing cybersecurity needs of U.S. election offices.

To learn more, visit [www.cisecurity.org](http://www.cisecurity.org).

The MS-ISAC is the focal point for cyber threat prevention, protection, response, and recovery for U.S. State, Local, Tribal, and Territorial (SLTT) governments. The Center for Internet Security's 24x7x365 Security Operations Center (CIS SOC) provides real-time network monitoring, early cyber threat warnings and advisories, vulnerability identification, and mitigation and incident response services to MS-ISAC members.

For more information, please visit <http://msisac.cisecurity.org>.



 [cisecurity.org](http://cisecurity.org)  
 [info@cisecurity.org](mailto:info@cisecurity.org)  
 518-266-3460  
 Center for Internet Security

 @CISecurity  
 TheCISecurity  
 cisecurity