TA14-017A: UDP-Based Amplification Attacks

Updated 02/27/2018 08:45 PM EST – Added information on Memcache-based reflection DDoS attacks

Original release date: January 17, 2014

Systems Affected

Certain application-layer protocols that rely on the User Datagram Protocol (UDP) have been identified as potential attack vectors. These include

- Domain Name System (DNS),
- Network Time Protocol (NTP),
- Connection-less Lightweight Directory Access Protocol (CLDAP),
- Character Generator Protocol (CharGEN),
- Simple Service Discovery Protocol (SSDP),
- BitTorrent,
- Simple Network Management Protocol version 2 (SNMPv2),
- Kad,
- Portmap/Remote Procedure Call (RPC),
- Quote of the Day (QOTD),
- Multicast Domain Name System (mDNS),
- Network Basic Input/Output System (NetBIOS),
- Quake Network Protocol,
- Steam Protocol,
- Routing Information Protocol version 1 (RIPv1),
- Lightweight Directory Access Protocol (LDAP),
- Trivial File Transfer Protocol (TFTP), and
- Memcache.

Overview

A distributed reflective denial-of-service (DRDoS) attack is a form of distributed denial-of-service (DDoS) that relies on publicly accessible UDP servers and bandwidth amplification factors (BAFs) to overwhelm a victim's system with UDP traffic.

Description

By design, UDP is a connection-less protocol that does not validate source Internet Protocol (IP) addresses. Unless the application-layer protocol uses countermeasures such as session initiation in Voice over Internet Protocol, an attacker can easily forge the IP packet datagram (a basic transfer unit associated with a packet-switched network) to include an arbitrary source IP address. [1] When many UDP packets have their source IP address forged to the victim IP address, the destination server (or amplifier) responds to the victim (instead of the attacker), creating a reflected denial-of-service (DoS) attack.

Certain commands to UDP protocols elicit responses that are much larger than the initial request. Previously, attackers were limited by the linear number of packets directly sent to the target to conduct a DoS attack; now a single packet can generate between 10 and 100 times the original bandwidth. This is called an amplification attack, and when combined with a reflective DoS attack on a large scale, using multiple amplifiers and targeting a single victim, DDoS attacks can be conducted with relative ease.

The potential effect of an amplification attack can be measured by BAF, which can be calculated as the number of UDP payload bytes that an amplifier sends to answer a request, compared to the number of UDP payload bytes of the request. [2] [3]

The following is a list of known protocols and their associated BAFs. US-CERT offers thanks to Christian Rossow for providing this information. For more information on BAFs, please see Christian's blog and associated research paper.

Protocol	Bandwidth Amplification Factor	Vulnerable Command
DNS	28 to 54	see: TA13-088A [4]
NTP	556.9	see: TA14-013A [5]
SNMPv2	6.3	GetBulk request
NetBIOS	3.8	Name resolution
SSDP	30.8	SEARCH request
CharGEN	358.8	Character generation request
QOTD	140.3	Quote request
BitTorrent	3.8	File search
Kad	16.3	Peer list exchange
Quake Network Protocol	63.9	Server info exchange
Steam Protocol	5.5	Server info exchange
Multicast DNS (mDNS)	2 to 10	Unicast query
RIPv1	131.24	Malformed request
Portmap (RPCbind)	7 to 28	Malformed request
LDAP	46 to 55	Malformed request [6]
CLDAP [7]	56 to 70	—
TFTP [23]	60	—
Memcache	10,000 to 51,000	—

In March 2015, the CERT Coordination Center of the Software Engineering Institute issued Vulnerability Note VU#550620 describing the use of mDNS in DRDoS attacks. Attackers can leverage mDNS by sending more information than can be handled by the device, thereby causing a DoS condition. [8]

In July 2015, Akamai Technologies' Prolexic Security Engineering and Research Team (PLXsert) issued a threat advisory describing a surge in DRDoS attacks using RIPv1. Malicious actors are leveraging the behavior of RIPv1 for DDoS reflection through specially crafted request queries. [9]

In August 2015, Level 3 Threat Research Labs reported a new form of DRDoS attack that uses portmap. Attackers are leveraging the behavior of the portmap service through spoofed requests to flood a victim's network with UDP traffic. [10]

In October 2016, Corero Network Security reported a new DDoS amplification attack exploiting LDAP directory services servers against its customers. [11]

In November 2017, Netlab 360 reported that CLDAP is now the third most common DRDoS attack, behind DNS and NTP attacks. [12]

In February 2018, SENKI reported an increase in Memcache-based reflection DDoS attacks (via UDP/TCP port 11211) with an unprecedented amplification factor. [24]

Impact

Attackers can use the bandwidth and relative trust of large servers that provide the UDP protocols provided in this alert to flood victims with unwanted traffic and create a DDoS attack.

Solution

Detection

Detection of DRDoS attacks is not easy because of their use of large, trusted servers that provide UDP services. Network operators of these exploitable services may apply traditional DoS mitigation techniques. To detect a DRDoS attack, watch out for abnormally large responses to a particular IP address, which may indicate that an attacker is using the service.

There are a few things victims of DRDoS attacks can do to detect such activity and respond:

1. Detect and alert large UDP packets to higher order ports.
2. Detect and alert on any non-stateful UDP packets. (A simple Snort example is below. The approach will need to be customized to each environment with a whitelist and known services.

   Simple Snort rule example for stateless UDP check

   ```
   var HOME_NET [10.10.10.20] preprocessor stream5_global: track_ip yes, track_tcp yes,track_udp
   no,max_tcp 262144, max_udp 131072 preprocessor stream5_ip: timeout 180 preprocessor stream
   use_static_footprint_sizes preprocessor stream5_udp: timeout 180, ignore_any_rules alert udp H
   (msg:"UDP Session start"; flowbits:set,logged_in; flowbits:noalert; sid: 1001;) alert udp any any ->
   (msg:"UDP Stateless"; flowbits:isnotset,logged_in; sid: 1002)
   ```

3. Upstream providers should maintain updated contacts and methods with downstream customers to send alerts by network.

In general, network and server administrators for Internet service providers (ISPs) should use the following best practices to avoid becoming amplifier nodes:

1. Use network flow to detect spoofed packets. (See the Mitigation section below for information on verifying spoofed traffic before blocking that traffic.)

2. Use network flow or other summarized network data to monitor for an unusual number of requests to at-risk UDP services.
3. Use network flow to detect service anomalies (e.g., bytes-per-packet and packets-per-second anomalies).

Mitigation

The following steps can help mitigate a DRDoS attack:

1. Use stateful UDP inspections—such as reflexive access control lists—to reduce the impact to critical services on border firewalls or border routers. [13]
2. Use a Border Gateway Protocol (BGP) to create a Remotely Triggered Blackhole, preferably in coordination with upstream providers or ISPs. [14]
3. Maintain a list of primary upstream provider emergency contacts to coordinate responses to attacks. Upstream providers should conduct mitigation in coordination with downstream customers.

In general, ISP network and server administrators should use the following best practices to avoid becoming amplifier nodes:

1. Regularly update software and configurations to deny or limit abuse (e.g., DNS response rate limit). [15] [16] [17]
2. Disable and remove unwanted services, or deny access to local services over the Internet.
3. Use UDP-based protocols—e.g., quality of service (QoS) on switching and routing devices—to enable network-based rate-limiting to legitimate services provided over the Internet.
4. Work with Customer Provider Edge manufacturers for secure configuration and software. [18]

As a service provider, to avoid any misuse of Internet resources:

1. Use ingress filtering to block spoofed packets (See the Spoofer Project [19], and IETF BCP 38 and BCP 84 guidelines). [20]
2. Use traffic shaping on UDP service requests to ensure repeated access to over-the-Internet resources is not abusive. [21] [22]

References

- [1] SIP: Session Initiation Protocol
- [2] Amplification Hell: Abusing Network Protocols for DDoS (link is external)
- [3] Amplification Hell: Revisiting Network Protocols for DDoS Abuse (link is external)
- [4] DNS Amplification Attacks

- [5] NTP Amplification Attacks Using CVE-2013-5211
- [6] Open LDAP Scanning Project
- [7] CLDAP Reflection DDoS (link is external)
- [8] VU#550620: Multicast DNS (mDNS) implementations may respond to unicast queries originating outside the local link
- [9] RIPv1 Reflection DDoS [Medium Risk] (link is external)
- [10] A New DDoS Reflection Attack: Portmapper; An Early Warning to the Industry (link is external)
- [11] Corero Warns of Powerful New DDoS Attack Vector (link is external)
- [12] CLDAP is Now the No.3 Reflection Amplified DDoS Attack Vector, Surpassing SSDP and CharGen (link is external)
- [13] Configuring IP Session Filtering (Reflexive Access Lists) (link is external)
- [14] Remotely-Triggered Black Hole (RTBH) Routing (link is external)
- [15] A Quick Introduction to Response Rate Limiting
- [16] Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing
- [17] Ingress Filtering for Multihomed Networks
- [18] Abuse of Customer Premise Equipment and Recommended Actions (link is external)
- [19] The Spoofer Project
- [20] Abuse of Customer Premise Equipment and Recommended Actions (link is external)
- [21] An Architecture for Differentiated Services
- [22] New Terminology and Clarifications for Diffserv
- [23] TFTP DDoS Amplification Attack (link is external)
- [24] Memcached on Port 11211 UDP & TCP Being Exploited

Revisions

- February 9, 2014 – Initial Release
- March 7, 2014 – Updated page to include research links
- July 13, 2015 – Added RIPv1 as an attack vector
- August 19, 2015 – Added Multicast DNS (mDNS) and Portmap (RPCbind) as attack vectors

- April 13, 2016 – Updated detection and mitigation information
- November 4, 2016 – Updated for LDAP attack vector
- December 4, 2017 – Added information on CLDAP as an attack vector
- December 6, 2017 – Added information on TFTP as an attack vector
- February 27, 2018 – Added information on Memcache as an attack vector