Preparing for denial-of-service attacks with Talos Incident Response

blog.talosintelligence.com/2022/03/preparing-for-denial-of-service-attacks.html

By Yuri Kramarz.

Over the years, several extorsion-style and politically motivated denial-of-service attacks increased and still pose a threat to businesses and organizations of any size that can find themselves in the crosshairs of various malicious campaigns.

A detailed preparation plan is needed to handle attacks that might come in various formats and over different protocols aiming at grinding existing infrastructure to a halt.

Understanding potential effects ahead of an attack can lead to reduced response times, which in turn would benefit business operations.

Executive Overview

Over the last few years, the world has seen several attacks aiming at temporarily disabling the infrastructure of the target companies or entire governments. These attacks, known as denial-of-service (DoS), or distributed denial-of-service (DDoS) use multiple network resources such as botnets or compromised systems to produce a staggering amount of network traffic against the target environment causing it to become unavailable.

To create such a disruption, adversaries use multiple resources to amplify the network traffic and conceal the identity of an attacker, thus complicating mitigation efforts. The end effect of a DoS or DDoS attack could vary from simply disrupting online services, to taking entire commercial businesses offline. While nothing can change the nature of adversaries out there, and their use of different tactics, tools and procedures (TTPs) against target companies, preparation for various types of malicious activities can reduce impact and recovery time thus, greatly reducing the cost of cyber incidents.

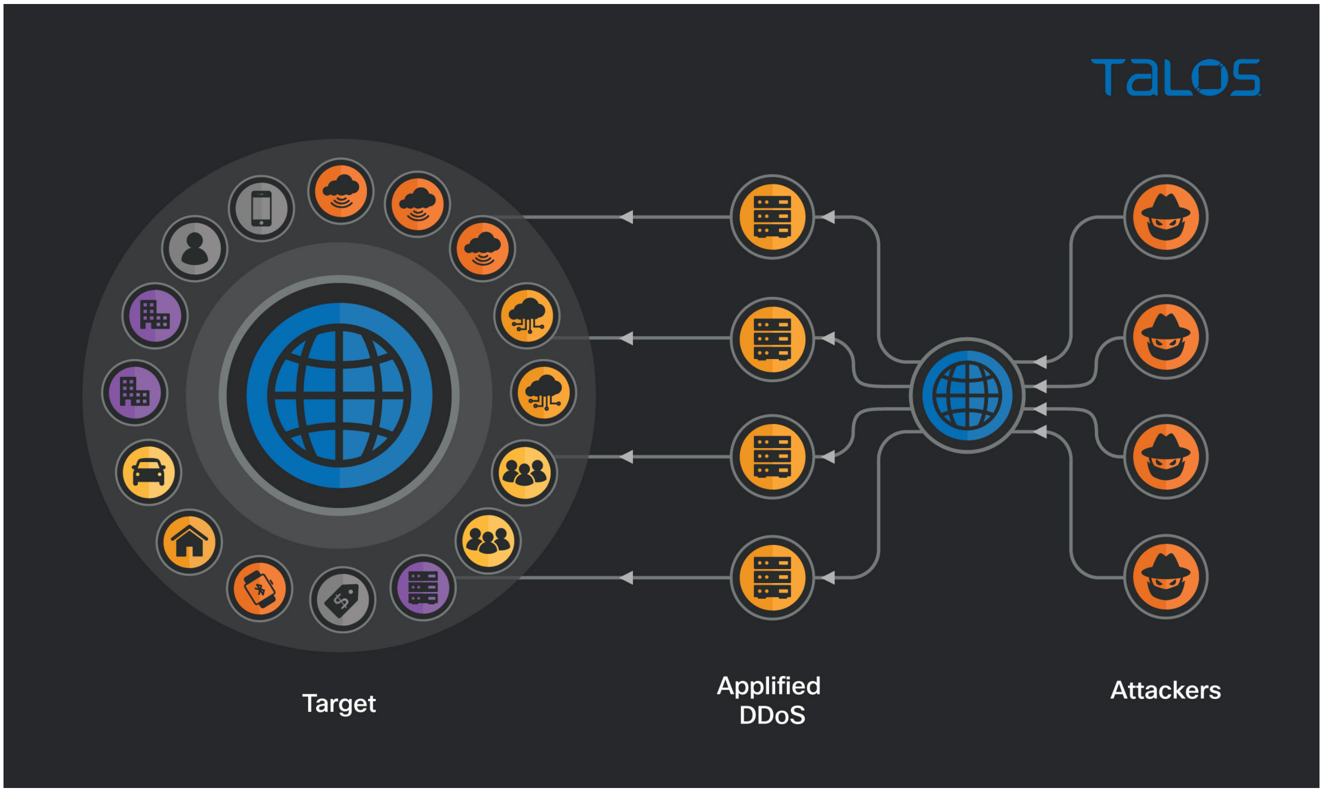

Technical overview

Over the years, security professionals have observed the bandwidth of DoS attacks increasing, as do their effects. In 2004, DDoS attacks used 8 Gbps on average. In 2015, these attacks were already around 400 Gbps and up to 2.3 Tbps by 2017.

This increase is mainly due to the use of amplification effects using attacks such as UDP fragmentation or DNS reflection, which are difficult to detect and contain without specialized tools. Another hard-to-identify vector observed in the wild corresponds to application layer attackers, where headless browsers or Slowloris-type attacks consume resources over the life of the requests to the point where, eventually, a server runs out of available sockets to create a new connection pool and is no longer accessible. Likewise, large botnets such as Mirai, use compromised IoT devices, making them capable of generating a significant amount of traffic in short amounts of time. Instead of using reflection or protocol misconfigurations, botnets generate traffic from thousands of compromised systems en masse.

Types of distributed denial-of-service attacks

There are three types of DDoS attacks, and each might require different procedures. These attacks are measured in bits per second.

Volumetric DDoS

Volumetric DDoS include the use of UDP, ICMP or other protocols that typically do not require handshaking. In this attack, the adversary aims to saturate all available bandwidth on the target network. The magnitude of this attack can be practically unlimited since botnets or other zombie computers send large amounts of traffic from around the globe. Since there are no handshakes associated with these protocols, it’s very easy to keep sending out packets toward the target. Amplification attacks also come under the same category, as they take advantage of weaknesses in protocols such as Network Time Protocol (NTP) or Domain Name System (DNS) queries to amplify requests that can then be multiplied in size.

Protocol DDoS

Protocol DDoS attacks rely on protocol weaknesses, which can be exploited to perform attacks such as SYN Floods, Smurf or fragmented packet attacks. These types of attacks consume processing power from networking devices such as load balancers or firewalls, thus making it nonresponsive to other entities due to use of the FIFO (first-in, first-out) principle applied for most devices processing network packets before actual servers serving the application and other resources. Typically, layers 3 and 4 are targeted in protocol DDoS attacks.

Application DDoS

Application DoS is caused by the submission of large POST requests, GET floods, or other application interaction techniques to overwhelm the application at layer 7 of TCP/IP stack. While possibly the hardest to execute, software programs (e.g., Slowloris) have been developed to perform these types of DDoS attacks. This type of attack sends data formatted in a specific application protocol (e.g., HTTP, SMTP, etc.) so it passes through the initial filtering on the boundary devices and reaches processing applications. As applications are often connected to a database, this attack may affect back-end systems, too.

Preparation for defence

Planning for DDoS attacks is crucial to Business Continuity Planning (BCP). Organizations need to assess its continuity planning in a formal risk assessment procedure that reflects the true state of the business’s infrastructure and services. The consequences of DDoS can be only understood and appreciated when the true attack surface of available services (i.e., network, connectivity, bandwidth, and devices) is known by the business.

As adversaries can shift their targets rapidly through botnets (i.e., Mirai), it's important to prepare for various scenarios when planning protections against DoS attacks.

Assess external attack surface by performing a formal asset discovery

Identify assets and the criticality of these assets as they relate to business continuity. This exercise should identify the systems that need to stay connected and any effects on the business should these systems face a DoS attack. The list of critical assets will be required to create an appropriate mitigation plan. Identify where the organization is most vulnerable from the attacker's point of view by reviewing network diagrams, business products sites, public IP addresses, exposed NAT addresses, database lists, system hosting details, IPAM and other data sources.

Assess the impact of the DoS attack on back-end systems such as databases, routers, and switches

Draw a dependency tree and data flow mapping between exposed infrastructure and various back-end systems to understand a worst-case scenario when under attack. This exercise should help you to understand the 'what-if' scenario before a full-scale attack is experienced by an organization.

Identify any systems that are a single point of failure

Based on the previous steps of identifying dependencies and dataflow mappings, ensure that any systems identified as a single point of failure can be brought up from backup whenever needed. If systems go down, a quick restore procedure can bring them up again. Test backup and restore procedures on a regular basis as part of at least annual BCP exercises. Where feasible, create a long-term strategic plan to migrate systems deemed as a single point of failure to a more resilient state.

Review your incident response (IR) plan

Document a solid business continuity plan that can be triggered when under attack and follow additional guidance in the incident response plan. Ensure that employees can use external systems and enable alternative routes into the network (e.g., a separate VPN with multi-factor authentication (MFA)) and notify them about attacks via out-of-bound methods such as text messages. Also, you'll want to document a process for notifying law enforcement or the local Computer Emergency Response Team (CERT) about the attack in progress.

Ensure that ISP contact details are known

An organization’s ISP will be crucial to mitigating some of the volumetric attacks when they appear. Communication with the ISP should be part of the IR plan by ensuring the ISP can contact appropriate technical leads for warning signs. Technical leads must understand the measures the ISP can perform to assist or mitigate a DDoS attack.

High-performance DDoS devices should be deployed externally to auto-mitigate DDoS attacks.

Establish contracts and service-level agreements (SLAs) for DDoS protection with companies who have the appropriate infrastructure and services to assist in the event of an attack. Cisco Secure Firewall appliances easily integrate with Radware, a Cisco Partner, to provide this for end customers. Business entities providing DDoS mitigation or protection services typically have different capabilities to absorb the DDoS traffic, which protects their customer networked environment from being saturated or forced offline.

Harden external devices, including operating system and application stacks

Ensure that applications and other components are stress-tested to handle exceptional traffic in a secure manner. Organizations need to focus on supporting libraries being patched and, where possible various hardening options deployed for underlying software.

Reduce “noise” in your environment through proper device configuration.

As system administrators and security personnel fight, an increasing number of automatic alerts are produced by an ever-growing number of devices in their environment. To avoid unnecessary noise in the environment, invest time in the hardened configuration with alert defenders with high fidelity alerts. If needed, involve the vendors and external service providers in fine-tuning the configuration. Revise the default settings and adjust various configurations to ensure that low-fidelity alerts are handled automatically.

Use Content Delivery Networks (CDNs) for static content.

Publish external content via CDNs, as these usually have DDoS protection enabled via services such as Cloudflare, Akamai or others.

Ensure that the organization is ready for an incident

By executing IR readiness assessments, tabletop and other exercises aimed at identification and preparation of appropriate procedures for when an incident does happen, your organization and its stakeholders can test its resilience before a real-world attack happens, and create awareness about their role in the organizational response.

Mitigation during a DDoS attack

While only a limited amount of mitigation techniques are possible without the involvement of specialized providers or tools, here are a few suggestions to reduce any unnecessary risk of being affected by DDoS attacks:

- Note that even with IPS in place, a large amount of traffic can overwhelm security devices. But it may be possible for defenders to block the attack at the external firewall, possibly with the added protection of an intrusion prevention system (IPS) like SNORT®. This is especially important for Layer 7 DDoS attacks. Reflective attacks will not be stopped by IPS, so other steps might be needed if this type of attack is executed against an organization.
- Organizations could also add on auto-scale resources where possible, especially when frontend servers, databases, or applications are cloud-based to increase CPU, memory and/or bandwidth.
- Work with ISP to redirect Volumetric DDoS to blackhole infrastructure, thus mitigating traffic that exceeds the organization's bandwidth.
- Use reverse path forwarding (RPF) on external routers to ensure that IP address spoofing is mitigated.
- Deny geolocation connections originating from countries where an organization does not have any business.
- Place limits on the traffic when feasible by configuring burst size and traffic priority on individual packet types. Where possible, deploy global ACL that can help to de-prioritize DDoS traffic to boundary devices if these are not overwhelmed by traffic coming through.
- Apply limits for ICMP, SYN packet rates and DNS TTL for exposed systems.
- Get a list of attacking IPs from logs and sort by traffic volume. Attempt to profile the traffic based on sample packet capture from perimeter devices. As the number of attackers might be quite substantial this might not always be feasible, however.
- Notify the executive board or CISO about the attacks and, if customers are impacted, issue a public statement describing the attack to inform them to ensure that the switchboard or IT support are not flooded with requests to assist.