CERT-EU

DDoS Overview and Response Guide

CERT-EU Team
ver. **2.0**
2/06/2024

Contents

1 Introduction

A *distributed denial-of-service* attack (or *DDoS* attack) is a malicious attempt using multiple systems to make computer or network resources unavailable to their intended users, usually by interrupting or suspending services connected to the Internet[1].

The specific methods and types of DDoS attacks change and evolve over time. The same applies to the victim targeting and motivation of the attackers. In this guidance, we focus on the situation in 2024 – based on experiences gathered when drafting earlier versions of this guidance as well as the most recent statistics of DDoS attacks from Q1 2024.

One of the most constant changes in DDoS attacks over the years has been their scale. Every year, systematically, the attacks are getting larger and more powerful. On 21 October 2016, a series of DDoS attacks against Dyn DNS (a DNS provider used by many important Internet companies) impacted the availability of a number of sites concentrated in the north-east of the United States and, later on, other areas of the USA. Impacted sites included, among others: PayPal, Twitter, Reddit, GitHub, Amazon, Netflix, Spotify, and RuneScape. Also, earlier that month, the attacks against the *Krebs on Security* blog[2] and the French Internet service and hosting provider OVH reached 620 Gbps and 1.2 Tbps respectively[3].

Botnets exploiting vulnerabilities in hundreds of thousands of Internet-connected devices, such as cameras, DVRs, and home DSL routers, were behind an important part of these massive attacks. The fact that security was not one of the design targets for these devices was long considered a potential threat by the security community, which was often neglected by the manufacturers.

Since the notable wave of DDoS attacks in 2016, DDoS attacks have become even more powerful. According to the CloudFlare blog[4], in Q1 2024, CloudFlare alone mitigated 4.5 million DDoS attacks, which together included 10.5 trillion HTTP DDoS requests and 59 PB (petabytes) of DDoS traffic. One of these attacks – launched by a Mirai-variant botnet – reached 2 Tbps.

Strategies to mitigate DDoS need to be adopted. These should focus primarily on prevention but eventually also include designing multilayered defence strategies. Therefore, DDoS threats should be taken into account as part of Business Continuity Planning, along with issues such as site selection, power outages, and natural disasters.

In this document, CERT-EU has focused on procedures for securing IT infrastructure from threats against **availability**. The guidance is based on proven DDoS identification and mitigation methods that can effectively and efficiently respond to DDoS attacks.

1.1 Target audience

This document is aimed at general IT staff who have undertaken the responsibility of managing Internet-facing IT infrastructures and being prepared to respond to DDoS incidents. This document only provides high-level guidelines. Different approaches are possible and may be valid. This document should rather be seen as a guideline in the absence of more specific local policies and procedures related to this topic. It does not supersede any specific applicable policies or procedures, which should be followed if they exist.

In case of doubts or any additional questions about this document, do not hesitate to seek further advice and assistance from your respective authorities or the CERT-EU team.

[1] http://www.incapsula.com/ddos/ddos-attacks/denial-of-service
[2] https://krebsonsecurity.com/tag/ddos/
[3] https://www.flashpoint-intel.com/action-analysis-mirai-botnet-attacks-dyn/
[4] https://blog.cloudflare.com/ddos-threat-report-for-2024-q1

2 DDoS attack categories

There are four primary categories of DDoS attacks[5]. Understanding these categories is essential for devising comprehensive defence strategies against DDoS attacks. Each type of attack requires a tailored approach for effective mitigation, often involving a combination of network, protocol, and application-layer defences.

Volumetric attacks: Volumetric attacks are designed to consume the bandwidth of the target network or service. They typically involve overwhelming the target with a massive amount of data, often through methods such as UDP floods, ICMP floods, or DNS amplification. To do so, the attacker generates large volumes of traffic by exploiting the capacity of many devices across the Internet, creating a flood of data packets. This type of attack is measured in bits per second (bps).

Protocol attacks: Protocol attacks aim to exhaust server resources or intermediate communication equipment (such as firewalls and load balancers) by exploiting protocol weaknesses. Examples include SYN floods, Ping of Death, and Smurf attacks. These attacks focus on depleting state table resources or causing excessive computational load on protocol stack implementations, thus disrupting legitimate connections. They are typically measured in packets per second (pps).

Application layer attacks: Application layer attacks target the application layer (Layer 7) of the OSI model, aiming to disrupt specific functions or features of a website or online service. Examples include HTTP floods, Slowloris, and DNS query floods. These attacks often mimic legitimate user behaviour, making them harder to detect. They focus on exhausting the application resources, leading to denial of service for legitimate users. These attacks are typically measured in requests per second (rps).

Resource exhaustion attacks: Although often overlapping with the other categories, resource exhaustion attacks specifically target the exhaustion of server or application resources, such as CPU, memory, or disk I/O. This can include attacks like R-U-Dead-Yet (RUDY) and Sockstress. These attacks focus on maintaining connections or consuming application resources at a rate that the server cannot handle, leading to a slowdown or complete halt of service.

3 DDoS threat landscape

The governmental sector is increasingly becoming a prime target for DDoS attacks. Since the onset of Russia's war on Ukraine, the frequency and intensity of DDoS attacks on governmental entities have notably increased. Threat actors increasingly leverage multi-vector strategies, combining volumetric, protocol, and application-layer attacks to overwhelm defences. For example, in October 2023, CloudFlare observed a DDoS attack with a peak of 398 million requests per second. Protecting against DDoS attacks is crucial, as maintaining service availability is essential for the operational health of any organisation, ensuring business continuity. Additionally, downtime can severely damage a company's reputation and customer trust, making reputation management a critical concern.

DDoS attacks in the governmental sector are frequently components of hybrid attacks. These attacks are designed not only to disrupt services but also to manipulate public perception and media narratives. Hacktivist personas on social media platforms often amplify DDoS attacks by exaggerating their impact. This amplification can influence the public and media, making the attack seem more significant than it actually is. It is therefore valuable to prepare higher management for the possibility of media reports following a DDoS attack. These reports can

[5] http://www.incapsula.com/ddos/ddos-attacks/denial-of-service

inadvertently assist hacktivists in spreading their narrative, thus achieving their influence objectives. Therefore, timely internal communication about the occurrence and real impact of DDoS attacks can help manage expectations and mitigate the influence of external narratives.

DDoS attacks against governmental institutions are often triggered by symbolic events. Some common triggers include:

- National elections
- Imposition of international sanctions
- Public support for geopolitical events, such as support for Ukraine
- Military aid announcements to conflict zones

4 DDoS mititgation

As mentioned, DDoS is one of the risks to be addressed in the organisation's Business Continuity Plan (BCP). The organisation should start by assessing the likelihood of different scenarios and the business impact on the organisation. Only after understanding the consequences of a DDoS attack and its likelihood, can the accountable managers for the potentially affected service start and support the necessary actions and plans to reduce the risk to a level they can accept.

CERT-EU recommends a formal approach to risk assessment for those constituents who have already been attacked, as well as for those who have a high risk as a result of a preliminary estimation. To achieve this, CERT-EU recommends using the same risk assessment methodology that is used for the rest of IT security risk mitigation, but focusing on the **availability** of the information, instead of confidentiality or integrity.

Defending a site against a DDoS attack has both a fixed and a variable cost. The fixed costs come in the form of locations, servers, and engineering. They can also include externally provided services, such as hosting, Content Delivery Network (CDN) services, or specific DDoS protection services. The variable, or operational, costs include the bandwidth served and manpower needed to mitigate attacks for the time they are ongoing, as well as possible variable costs related to Cloud services, such as CDN or DDoS protection. From this perspective, mitigating DDoS is a business decision that should address what service should be still available under what kind of DDoS attack and for how long. It is also important to support the proper incident response plan with the proper budget allocations and to accept the residual risk. It is important to remember that starting by considering only (or mainly) technical aspects might be misleading.

Some additional resources related to preparation for DDoS attacks that are available online are presented in the table below.

Title	Comment
How to prevent DDoS attacks	CloudFlare guide
AWS Best Practices for DDoS Resiliency	AWS guide
Azure DDoS Protection documentation	Azure documentation
Understanding and responding to DDoS Attacks	CISA guide
DoS guidance	NCSC-UK
Gartner DDoS Mitigation Solutions	Gartner

Having this information in mind, CERT-EU proposes the following response plan.

4.1 Mitigation techniques

Effective DDoS mitigation requires a multi-layered approach, combining several techniques to protect against various types of attacks. By implementing these strategies, organisations can significantly reduce the impact of DDoS attacks and ensure the availability and reliability of their services. Here are some of the most common and efficient DDoS mitigation techniques:

- **Content Delivery Networks (CDNs)**: CDNs proxy content delivered from the back-end (or *origin*) servers using a distributed network of servers to reduce the load on the origin server. CDNs can deliver cached copies of static content (e.g., images, videos, scripts, etc.) from edge servers close to users. It is effective against large-scale volumetric attacks by distributing the load. Often CDNs will also *hide* origin servers' IP addresses by redirecting DNS records to the CDN's edge servers. This additionally isolates the organisation from direct attacks. CDN solutions often provide the ability to temporarily or permanentely enable a javascript challenge that will require a manual validation from the visitors before reaching the website. This is generally a good counter-measure against application layer DDoS attacks.
- **Web Application Firewalls (WAFs)**: WAFs protect web applications by filtering and monitoring HTTP/HTTPS requests. They use predefined rules to identify and block malicious traffic, and identify unusual patterns and behaviours that could indicate an attack. It is effective against application layer attacks, such as SQL injection and cross-site scripting (XSS). WAFs can be deployed at the CDN level or at the back-end servers level.
- **Load balancing**: Distributes incoming traffic across multiple servers to prevent any single server from becoming overwhelmed. Load balancers direct traffic to different servers based on load, health, and other factors. It is effective against both volumetric and application layer attacks by ensuring that traffic is evenly distributed. Again, load balancers can be deployed at the CDN level or at the back-end servers level.
- **Auto-scaling**: Auto-scaling is a cloud service feature that automatically adjusts the number of active servers or resources based on current demand. It ensures that applications maintain performance and availability despite traffic spikes. However, continuous scaling up during a sustained DDoS attack can lead to significant costs if not properly configured and limited.
- **Rate limiting**: This technique involves controlling the rate at which requests are processed by the servers to prevent overload. It works by limiting the number of requests a user can make in a certain timeframe. It can be applied at various levels, including web servers, APIs, and applications. It is effective against application layer (Layer 7) attacks such as HTTP floods but might impact the availability of resources to legitimate clients.
- **Rate-based packet filtering**: Filter packets based on predefined rate thresholds to block excessive traffic. It is effective against volumetric and protocol attacks by controlling the rate of incoming traffic. Again, it might impact the availability of resources to legitimate clients.
- **Upstream filtering**: ISPs and upstream providers can filter malicious traffic before it reaches the target network by implementing filtering rules to block attack traffic at a higher level. It is effective against large-scale volumetric attacks by stopping the traffic upstream before it reaches the target.
- **Traffic filtering and scrubbing**: Filtering and scrubbing involve cleaning the incoming traffic to separate legitimate requests from malicious ones. It works by rerouting traffic to scrubbing centres where malicious traffic is filtered out before reaching the target. The scrubbing centres will block traffic from known malicious IP addresses and analyse

the content of data packets to identify and block malicious traffic. It is effective against volumetric attacks and protocol attacks (e.g., SYN floods). However, the cost of scrubbing services usually depends on the volume, so it could be potentially high for bigger attacks.

4.2 DDoS response plan

4.2.1 Preparation

The preparation phase is critical in ensuring that an organisation is ready to effectively mitigate and respond to DDoS attacks. This phase involves establishing a strong foundation of policies, architectures, tools, and procedures, as well as training personnel to handle potential incidents.

Objective: Establish and maintain the necessary architectures, tools, policies, and procedures to mitigate and respond to DDoS attacks effectively.

1. Risk assessment

- **Identify critical assets**: Determine which systems, applications, and services are most critical to your operations and would be most impacted by a DDoS attack.
- **Vulnerability analysis**: Conduct regular vulnerability assessments to identify potential weaknesses in your network and application infrastructure that could be exploited in a DDoS attack.
- **Impact analysis**: Evaluate the potential impact of a DDoS attack on different parts of your infrastructure to prioritise resources and response strategies.

2. Incident response plan development

- **Roles and responsibilities**: Clearly define the roles and responsibilities of all team members involved in the incident response process, including IT staff, security personnel, and management.
- **Communication plan**: Establish internal and external communication protocols, including how to notify stakeholders and partners during an attack.

3. Training and awareness

- **Simulation exercises**: Perform regular DDoS attack simulations and tabletop exercises to test the effectiveness of your incident response plan and to identify areas for improvement.

4. Technology deployment

- **Protect applications**:
  - **Content Delivery Networks (CDNs)**: Use CDNs to distribute content geographically, reducing the load on your primary (origin) servers.
  - **DDoS protection services**: Subscribe to DDoS protection services offered by reputable providers. These services often include traffic scrubbing, rate limiting, and advanced filtering.
  - **Firewall and WAF configuration**: Properly configure firewalls and Web Application Firewalls (WAFs) to block malicious traffic and protect against application-layer attacks.
  - **Load balancers**: Use load balancers to distribute traffic evenly across multiple servers, preventing any single server from becoming overwhelmed.
  - **Redundancy and scalability**: Design your network to include redundancy and scalability to handle sudden spikes in traffic.
- **DDoS detection**:
  - **Monitoring tools**: Implement comprehensive network and application monitoring tools to detect unusual traffic patterns indicative of a DDoS attack.

5. Collaboration and Partnerships

- **ISP and upstream providers**: Establish relationships with ISPs and upstream providers to ensure they can assist in filtering attack traffic before it reaches your network.
- **Information sharing**: Stay informed about the latest DDoS attack vectors and mitigation techniques.

4.2.2 Identification

Objective: Detect and verify potential DDoS attacks promptly.

- **Anomaly detection**: Use monitoring tools to identify abnormal traffic patterns, such as unexpected spikes in bandwidth usage or an unusually high number of requests per second.
- **Traffic analysis**: Analyse network traffic logs to distinguish between legitimate traffic and potential attack vectors.
- **Alerting mechanisms**: Set up automated alerts for signs of potential DDoS attacks, enabling a swift response.

4.2.3 Containment

Objective: Limit the impact of the DDoS attack and prevent it from affecting additional systems.

- **CDN provided javascript challenge**: Enable the CDN javascript challenge at the CDN level.
- **Traffic filtering**: Apply rate limiting, blacklisting, and geofencing to filter out malicious traffic.
- **Upstream mitigation**: Collaborate with CDN providers, ISPs, and upstream providers to filter attack traffic before it reaches your network.
- **Traffic diversion**: Utilise scrubbing centres or redirect traffic through DDoS mitigation services to clean the incoming traffic.
- **Access control**: Temporarily disable non-essential services and restrict access to critical systems to reduce the attack surface.

4.2.4 Eradication

Objective: Identify and eliminate the source of the DDoS attack, if possible.

- **Forensic analysis**: Conduct a detailed analysis of the attack to identify its source and nature.
- **Vulnerability remediation**: Address any vulnerabilities that may have been exploited during the attack.

4.2.5 Recovery

Objective: Restore normal operations and ensure all systems are fully functional.

- **System restoration**: Gradually reintroduce services, starting with the most critical, to ensure stability.
- **Performance monitoring**: Continuously monitor system performance to detect any lingering issues or secondary attacks.
- **User communication**: Inform stakeholders and users about the attack, the steps taken to mitigate it, and the current status of services.

4.2.6 Lessons learned

Objective: Review the incident to improve future response and mitigation strategies.

- **Post-incident analysis**: Conduct a thorough post-mortem analysis to evaluate the effectiveness of the response and identify areas for improvement.
- **Documentation update**: Update the IR plan, policies, and procedures based on the findings from the incident analysis.
- **Training update**: Incorporate lessons learned into ongoing training programmes to enhance the organisation's readiness for future attacks.

5 Conclusion

Some time ago, DDoS attacks targeting public services might have appeared to be mostly linked to *hacktivism*. Nowadays, it is not possible to find any important political movement or campaign without an Internet presence. DDoS attacks are not merely technical disruptions but strategic tools used during geopolitical tensions to cause chaos and instil uncertainty.

The new DDoS tools and methods can be easily used for criminal purposes and financial gain, but also potentially to **achieve political means** for certain groups, parties, or even **nation-states**. These possibilities have to be kept in mind when evaluating and deciding how to prepare against potential DDoS attacks.

Consequently, the main ideas that CERT-EU would offer for consideration in order to face such issues are:

- Institutions should establish their availability requirements clearly. It might have different implications for critical services that must be available (even with legal implications), such as publications in the Official Diary or call-for-tenders, than for resources that might support and explain public policies or political decisions.
- It is very important to know in advance the weakest links and bottlenecks that might threaten these availability requirements in the case of a DDoS attack. To do that, a comprehensive risk assessment is highly recommended.
- Taking into account the latest attack strength, it is highly recommended to use CDNs and cloud-based DDoS protection services. It is also beneficial to hire specialist help or consultancy services to address demanding availability requirements.

Finally, in the case of a DDoS attack, reporting the incident and investigating various aspects, such as the threat actors involved or the techniques used, can help improve global security on the Internet.

TLP Definition

TLP	Disclosure	Message
RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed.

TLP	Disclosure	Message
AMBER	Limited disclosure, recipients can only spread this on a need-to-know basis within their organisation and its clients.	Recipients may share TLP:AMBER information only with members of their own organisation.
AMBER+STRICT	Limited disclosure, recipients can only spread this on a need-to-know basis within their organisation only.	Recipients may share TLP:AMBER+STRICT information only with members of their own organisation.
GREEN	Limited disclosure, restricted to the community.	Subject to standard copyright rules, TLP:GREEN information may be distributed with peers and partner organisations within their sector or community, but not via publicly accessible channels.
CLEAR	Disclosure is not limited.	TLP:CLEAR Recipients can spread this to the world, there is no limit on disclosure.