



Protecting your business against
financial and reputational losses with
Kaspersky DDoS Protection



A Distributed Denial of Service (DDoS) attack is one of the most popular weapons in the cybercriminals' arsenal. It aims to make information systems such as websites or databases impossible for regular users to access normally. There can be different motives behind launching DDoS attacks, ranging from cyber-hooliganism to dirty competition practices or even extortion.

The modern DDoS industry is a multi-layered structure. It includes people who commission attacks, the botnet creators who make their resources available, intermediaries who arrange the attacks and talk to the clients; and the people who arrange for payments for all the services provided. Any network node available in the Internet may become a target, be it a specific server, a network device or a disused address in the victim sub-network.

There are two common scenarios for conducting DDoS attacks: sending requests directly to the attacked resource from a large number of bots, or launching a DDoS amplification attack through publicly available servers containing software vulnerabilities. In the first scenario, cybercriminals turn a multitude of computers into remotely controlled "zombies" which then follow the master's command and simultaneously send requests to the victim computing system (conduct a "distributed attack"). Sometimes, a group of users is recruited by hackers, provided with special software designed to conduct DDoS attacks and given orders to attack a target.

Under the second scenario involving an amplification attack, servers leased out from a data center can be used instead of bots. Public servers with vulnerable software are typically used for enhancement. Today, either DNS (domain name system) servers or NTP (network time protocol) servers can be used. An attack is amplified by spoofing return IP addresses and sending a short request to a server that requires a much longer response. The received response is sent to the spoofed IP address which belongs to the victim.




Figure 1. Flow diagram of most popular versions of DDoS attacks

There is another factor that makes the situation even more dangerous. Because there is so much malware out there, and cybercriminals have created so many botnets, almost anybody can launch this kind of attack. Cybercriminals advertise their services saying that anybody can take down a specified site for just \$50 a day. The payments are typically made in cryptocurrency, so it is almost impossible to track down the orders through cash flows.

Affordable prices mean that any online resource can be targeted in a DDoS attack. It's not something limited to the Internet resources of large and famous organizations. It is more difficult to cause damage to web-resources owned by large companies, but if they are made unavailable, the cost of that downtime will be much greater. Apart from the direct losses resulting from missed business opportunities (such as electronic sales), companies can face fines for defaulting on their obligations or expenses relating to extra measures to protect themselves from further attack. Last but not least, company's reputation may be damaged, causing it to lose existing or future clients.

The total cost depends on the size of the business, the industry segment it serves and the type of service under attack. According to calculations by the analytical company IDC, one hour downtime of an online service can cost a company \$10,000 – \$50,000.

Methods of countering DDoS attacks

There are dozens of companies on the market that provide services to protect against DDoS attacks. Some install appliances in the client's information infrastructure, some use capabilities within ISP providers and other channel traffic through dedicated cleaning centers. However, all these approaches follow the same principle: junk traffic, i.e. traffic created by cybercriminals, is filtered out.

Installing filtering equipment on the client's side is considered to be the least effective method. Firstly, it requires specially trained personnel within the company to service the equipment and adjust its operation, creating extra costs. Secondly, it is only effective against attacks on the service, and does nothing to prevent attacks choking the Internet channel. A working service is of no use if it cannot be accessed from the net. As amplified DDoS attacks become more popular it has become much easier to overload a connection channel.

Having the provider filter the traffic is more reliable as there is a broader internet channel and it is much harder to clog it up. On the other hand, providers do not specialize in security services and only filter out the most obvious junk traffic, overlooking subtler attacks. A careful analysis of an attack and a prompt response require the appropriate expertise and experience. Besides, this kind of protection makes the client dependent on a specific provider and creates difficulties if the client needs to use a backup data channel or to change its provider.

As a result, specialized processing centers implementing a combination of various traffic filtration methods should be considered the most effective way to neutralize DDoS-attacks.

Kaspersky DDoS Protection

Kaspersky DDoS Protection is a solution that protects against all types of DDoS attacks by using a distributed infrastructure of data cleaning centers. The solution combines different methods, including traffic filtration on the provider side, installation of a remotely controlled appliance to analyze traffic next to the client's infrastructure, and the use of specialized cleaning centers with flexible filters. In addition the solution's work is constantly monitored by Kaspersky Lab's experts, so the onset of any attack can be detected as soon as possible, and filters can be modified as required.




Figure 2. Kaspersky DDoS Protection: Operation Diagram

Kaspersky Lab's arsenal

For more than a decade Kaspersky Lab has successfully dealt with a wide range of online threats. Over that time Kaspersky Lab's analysts have acquired a unique level of expertise, including a detailed understanding of how DDoS attacks work. The company's experts constantly watch the newest developments taking place on the Internet, analyze the latest methods of conducting cyber-attacks, and improve our existing protection tools. With this expertise at hand, it is possible to detect a DDoS attack as soon as it is launched and before it floods the target web resource.

The second element in Kaspersky's DDoS Protection technology is a sensor installed next to the client's IT infrastructure. The sensor is a piece of software running under the Ubuntu operating system and requiring a standard x86 server. It analyzes the types of protocols used, the number of bytes and data packets sent, the client's behavior on the web-site, i.e. the metadata, or information about the sent data. It does not redirect traffic anywhere, modify it or analyze the content of any messages. The statistics are then delivered to the cloud-based Kaspersky DDoS Protection infrastructure, in which a statistics-based profile is created for each client based on the collected metadata. In effect these profiles are records of typical information exchange patterns for each client. Changes in typical times of use are recorded. Later on, traffic is analyzed; any time the traffic behavior is different from the statistics-based profile it may be indicative of an attack.

The keystone of Kaspersky DDoS Protection is its cleaning centers. These are located on the main internet backbone lines, in places like Frankfurt and Amsterdam. Kaspersky Lab simultaneously uses several cleaning centers, so it can divide or redirect the traffic that needs to be cleaned. The processing centers are united into a common cloud-based information infrastructure and the data is contained without those boundaries. For example, the web traffic of European clients does not leave European territory.

Another key way of controlling DDoS-traffic is to filter it on the provider side. The ISP does not just supply an Internet channel, it can also enter a technology partnership with Kaspersky Lab. Thus, Kaspersky DDoS Protection can cut off the most obvious junk traffic, used in the majority of DDoS-attacks, as close to its point of origin as possible. This prevents the streams from merging into a single powerful attack and eases the burden on the cleaning centers, which are free to handle more sophisticated junk traffic.

Traffic redirection tools

For the security solution to work effectively, the first key requirement is to set up a connection channel between the cleaning centers and the client's IT infrastructure. In Kaspersky DDoS Protection, these channels are arranged according to the Generic Routing Encapsulation protocol. They are used to create a virtual tunnel between the cleaning center and the client's network equipment, through which the cleaned traffic is delivered to the client.

The actual traffic redirection can be done using one of two methods: by announcing the client's subnet using a BGP dynamic routing protocol, or by modifying the DNS record by introducing the URL of the cleaning center. The first method is preferable as it can redirect traffic much faster and protect against attacks that directly target a specific IP address. However this method needs the client to have an address range that is independent of the provider such as a block of IP addresses provided by a regional Internet registrar.

When it comes to the actual redirection procedure, there is little difference between the two methods. If the first method is used, then the BPG routers on the client's side and at the cleaning center establish a permanent connection via the virtual tunnel; in case of attack, a new route from the cleaning center to the client is created. When the second method is used, the client is assigned an IP address from the cleaning center's address pool. If an attack begins, the client replaces the IP address in the DNS A-record with the IP address assigned by the cleaning center. After this all the traffic arriving at the client's address will be sent to the cleaning center first. However, to stop the attack on the old IP address continuing, the provider has to block all incoming traffic except data coming from the cleaning center.

How it works

In normal circumstances, all traffic from the Internet goes directly to the client. The protective actions begin as soon as a signal from the sensor arrives. In some cases, Kaspersky Lab's analysts know about an attack as soon as it starts, and inform the client. In this case preventive measures can be taken in advance. The on-duty DDoS expert at Kaspersky Lab receives a signal that traffic arriving to the client does not match the statistical profile. If the attack is confirmed, then the client is notified of the attack, and should give the order to redirect the traffic to the cleaning centers (in some cases, there may be an agreement with the client that the redirection starts automatically).

As soon as Kaspersky Lab's technologies determine the type of the attack, specific cleaning rules are applied for this type of attack and the specific web resource. Some of the rules, designed to treat the crudest type of attacks, are communicated to the provider's infrastructure and are applied on routers owned by the provider. The remaining traffic is delivered to the cleaning center's servers and filtered according to a number of characteristic signs, such as IP addresses, geographical data, information from the HTTP headers, the correctness of protocols and exchange of SYN packets, etc.

The sensor continues to monitor the traffic as it comes to the client. If it still shows signs of a DDoS attack, the sensor alerts the cleaning center, and the traffic undergoes deep behavior and signature analysis. With these methods, malicious traffic can be filtered out based on signatures, i.e. a specific type of traffic can be completely blocked, or IP addresses can be blocked based on specific observed criteria. This way, even the most sophisticated attacks are filtered, including an HTTP flood attack. These attacks involve imitations of a user visiting a web-site, but they actually are chaotic, unnaturally fast, and typically come from a regiment of zombie computers.

Kaspersky Lab’s experts monitor the entire process using a dedicated interface. If an attack is more complicated than usual or is atypical, the expert may step in, change the filtering rules and reorganize the processes. Clients can also watch how the solution performs and how the traffic behaves, using their own interface.




Figure 3. Screenshot of the client's interface

When the attack is over, traffic is directed back to the client’s servers. Kaspersky DDoS Protection reverts to standby mode, and the client receives a detailed report of the attack, including a detailed account of how it developed, graphs plotting measurable parameters, and the geographical distribution of the attack sources.

Advantages of Kaspersky Lab's approach

- Only redirecting traffic to Kaspersky Lab cleaning centers during an attack and filtering traffic on the provider's side helps significantly reduce the cost to the customer.
- Filtration rules are individually developed for each customer depending on the specific online services that need to be protected.
- Kaspersky Lab experts monitor the process and quickly adjust filtration rules when necessary.
- Close cooperation between Kaspersky DDoS Protection experts and Kaspersky Lab developers makes it possible to adapt the solution flexibly and rapidly in response to changing circumstances.
- To ensure the highest possible level of reliability, Kaspersky Lab only uses European equipment and service suppliers in European countries.
- Kaspersky Lab has accumulated a wealth of experience applying this technology in Russia, where it successfully protects leading financial institutions, commercial and government agencies, online shops, etc.