EBU

OPERATING EUROVISION AND EURORADIO

R 141

MITIGATION OF DISTRIBUTED DENIAL OF SERVICE (DDoS) ATTACKS

RECOMMENDATION

SOURCE: SP-MCS

Geneva
June 2015

[*] Page intentionally left blank. This document is paginated for two sided printing

Mitigation of Distributed Denial of Service (DDoS) attacks on media companies

EBU Committee	*First Issued*	*Revised*	*Re-issued*
SP-MCS	2015		

Keywords: Cybersecurity, DDoS, Botnets, IP, Cyberthreat, Cyberattack.

Recommendation

The EBU, considering that,

1. Media companies increasingly have services on the internet (conventional Websites, Blogs, Community-Sites and streaming services for Radio and TV programme distribution).

2. The Cyberthreat of DDoS-attacks is present and is especially increasing towards Media- and Broadcast-Organisations.

3. DDoS is one of the most common and frequently encountered type of Cyberattack made on media companies.

4. DDoS attacks are often correlated to the content of the broadcast Radio or TV programmes and have an economic or political background.

5. From the attackers view, DDoS attacks are easier and cheaper to conduct than most types of Cyberattack.

6. DDoS attacks usually put an excessive load on media companies' critical infrastructure, which results in the unavailability of their services (Websites, Radio and TV programmes).

Recommends that:

The following security measures be implemented at any media facility.

1. **Detective Security** in the form of continuous monitoring of capacity (system CPU load, inbound traffic load etc.), traffic type and any critical infrastructure and service (e.g. firewalls etc.) with a view to improve detection capabilities of cyberattack.

2. **Preventive Security** measures such as:

   a. Segmentation of internal to external networks and of any network containing critical broadcast systems.

   b. Scanning (preferably automated) and patching of potential DOS vulnerabilities in internet facing services.

   c. Load Balancing.

d. Having proper tools and procedures in place to respond to attacks.

e. Defining a DDoS protection agreement with your ISP.

3. **Corrective Security** measures (at least one, preferably more) such as:

a. DDoS protection services that allow traffic cleaning (may be external ISP or third party based).

b. Internal edge security gateways that protect every system reachable from the internet. The security gateway itself should include (D)DOS detection and protection capabilities.

c. Additional network based countermeasures may be considered, e.g.:

Blackholing.

Blocking attackers IP addresses.

Stop IP announcing.

DNS reconfiguration.

Isolation (disconnect from internet access) - as last resort.

Note: Security solutions offering DDoS protection usually contain both detective and corrective measures. In some solutions, preventive techniques may also be present.

[Informative annex overleaf]

Annex: Background information on DDoS attacks and their potential countermeasures

This annex contains high-level descriptions of all countermeasure techniques (detective, preventive and corrective).

1. Background

A DDoS (Distributed Denial of Service) attack is a network based attempt to make a website, a service or a complete infrastructure unavailable, typically by simultaneously attacking a victim from several compromised systems. The attack types vary; the most common ones are described in § 3.

Media companies are increasingly reliant on IP based production and frequently offer their services on the internet, which exposes them to potential attackers.

DDoS attacks are becoming technically easier and cheaper to perform; they are becoming a major threat.

2. Potential business impact.

For media companies that offer internet based services (e.g. websites, on demand services, hybrid TV technologies,) DDoS attacks are a major threat, as, depending on the architecture of the infrastructure, a successful DDoS attack can harm the attacked services as well as the underlying or connected infrastructures (e.g. internal networks) too.

As a worst case, a DDoS attack against a website could also lead to breakdowns of internal radio and TV equipment if these systems are not properly separated or if supplying services such as DNS, AD or DHCP.

Several EBU Members have already suffered a wide range of DDoS attacks that led to critical damage to their infrastructures and rendered their internet based services unavailable.

3. Type of DDoS Attacks (General)

DDoS comprises several attack methods that address the Network, Session or Application Layers. DDoS attacks can be roughly distinguished as being symmetric or asymmetric, depending on whether the load is generated and spread symmetrically on both ends of the attack or not.

With symmetric attacks the attackers need to generate the complete load with their own resources and transfer it to the victim. Botnets (a large number of compromised, internet connected systems working in concert) are often involved in this type of attack.

With asymmetric attacks the attacker leverages the asymmetric nature of some internet protocols that can trigger a large reply (quantity of data) from a small request (quantity of data). This causes a small load on the attacker side, but results in a huge load, with a failure or consumption of critical resources on the victim’s side.

Symmetric attacks:

Dependent on the current broadcast content, media companies are highly susceptible victims for this form of DDoS attack, because of the small (financial and technical) effort needed to launch them. "Hacktivism-Tools" such as the "Low Orbit Ion Canon" and others are available to the public and facilitate the attack; otherwise, many websites offer high volume, botnet-based DDoS attacks for a relative small amount of money (60 Gbit/s for 25$/h).

Assymetric attacks:

a) Reflective/Amplifying attacks:

These kinds of attack make use of the fact that sending small or malformed requests to any target will result in the target replying to the sender with a much larger package than the original request. By spoofing the source address the reply is sent to the real target of the hacker. Usually these kinds of attack require a large number of systems that simultaneously perform requests with a common spoofed source address so that the replies will flood the victim's service or network. These systems are previously hacked systems that form a "botnet" that can be rented or used "as a service" for attacks.

b) SYN flood attacks:

With this still very popular attack the attacker sends a great amount of TCP SYN-Requests to a single IP-Target to fill up its limited connection table of half-open sessions, which is relatively easy to accomplish. Once this is achieved the affected host is unable to accept any new sessions.

c) Low bandwidth attacks:

These kinds of attack are based on the concept of consuming much or all resources on a webserver or application by starting a regular request, such as an HTTP POST, and then proceeding at an extremely slow rate (e.g. 1 byte/110 seconds) to complete the request. With this approach web server resources will be used up in an extremely short period of time and with very little effort on the attacker side. Known attacks of this type are Slowloris, SLOW post and RUDY (R-U-Dead-Yet?).

d) Fragmentation attacks:

Flooding target systems with many small IP fragments, zero-length TCP window packets or mangled IP fragments (teardrop) are also popular DoS attacks that increase the CPU consumption of a system. In this scenario, internal buffers and tables reach their storage limits, which will then lead to a breakdown of the system if no internal protection measures are in place in the target system.

e) IPv6 attacks:

With the emergence of IPv6 a whole plethora of attacks against their availability also emerged, including long-presumed-dead attacks such as Ping-of-Death or Landattack. See *THC IPv6 Attack Toolkit*.

f) Specialized attacks:

The strange or undefined behaviour of protocols or software in general is also a popular starting point for DDoS attacks. Because of the sheer wealth of attacks that are possible, only a few of the most popular in this category are given:

- **DNS NXDOMAIN**: Flooding a DNS Server with randomly-generated non-existent hostnames, which causes enormous workload on the DNS resolver making the whole server unavailable.
- **SSL Renegotiation DoS**: Causes a big load on the SSL-Server, when client-side initiated SSL

renegotiation is enabled.

- **HashDoS**: Extremely effective and clever attack. One especially crafted HTTP-POST can saturate a vulnerable webserver for hours. The problem lies within the way web frameworks such as PHP and many others generate and organize their internal dictionaries as hash tables. This attack is very hard to detect!
- **Apache Killer**: Vulnerable Apache versions consume most of the CPU by trying to calculate specially crafted byte range requests in HTTP 1.1.
- **HTTP Pipelining**: Misuse of this protocol features leads to complete saturation of the webserver. A few malicious clients using pipelining can completely fill up the server buffers with HTTP requests without the need to wait for any server response.
- **ReDoS**: Exploits the fact that most regular expression implementations may reach extreme situations that cause them to work very slowly (exponentially related to input size).

4. Potential Solutions

4.1 General Solutions

There are several options to mitigate a DDoS attack. Some of the proposed corrective security measures are able to successfully stop a DDoS attack from reaching the service it is aimed at. Unfortunately, legitimate packets are discarded along with the malicious traffic, so the service is unavailable to the outside world and the Denial of Service attack is successful.

If the goal is to keep the service available for legitimate users while the attack is ongoing, a system has to accurately detect an attack, distinguish legitimate traffic from DDoS traffic, and only forward legitimate traffic to the destination service. There are a couple of options for successfully excising DDoS traffic while still receiving legitimate traffic:

1. Hosted or 'as a Service' companies can activate their anti-DDoS services that are provided “externally” when necessary.

2. On-premises “internal” detection and mitigation using an appliance in the company’s own network that can detect and successfully mitigate DDoS attacks.

With the second option, there are more decisions to be made:

- **Detection using OSI layer 4 or layer 7 information**: To detect a DDoS attack, all the packets bound for the internal network need to be analysed. This can be done using Netflow data, which only gives up to OSI layer 4 information (*src*/*dst* IP and *src*/*dst* port), or using SPAN or inline data collection, which gives up to OSI layer 7 information (including, for instance, the HTTP commands in the packet).
- **In-band or out-of-band**: Some appliances need to be put in-band, which means that the appliance is physically on a link between two routers. Other appliances can be used out-of-band. These latter collect data using Flow or SPAN, and can divert the potentially bad traffic towards them.

All those options have advantages and drawbacks (‘pros’ and ‘cons’), and vary between network designs and solutions.

4.2. Examples of proven solutions for media companies

4.2.1 Solution 1 - External ISP based solution

Many ISPs (Internet Service Providers) provide DDoS protection or detection functionalities as add-ons to their internet access service. The following high-level solution is based on a typical DDoS protection platform that consists of two main systems:

1) The Netflow Collector

This monitors the inbound and outbound traffic based on Netflow data. By adjusting to a fitting sample rate the collector can create a very typical baseline traffic profile that it uses for creating alerts if differing traffic is then detected.

2) The Threat Management System (TMS)

The TMS makes use of several filtering technologies to "wash" the traffic. (E.g. excising a specific types of ICMP message or malformed http packages).

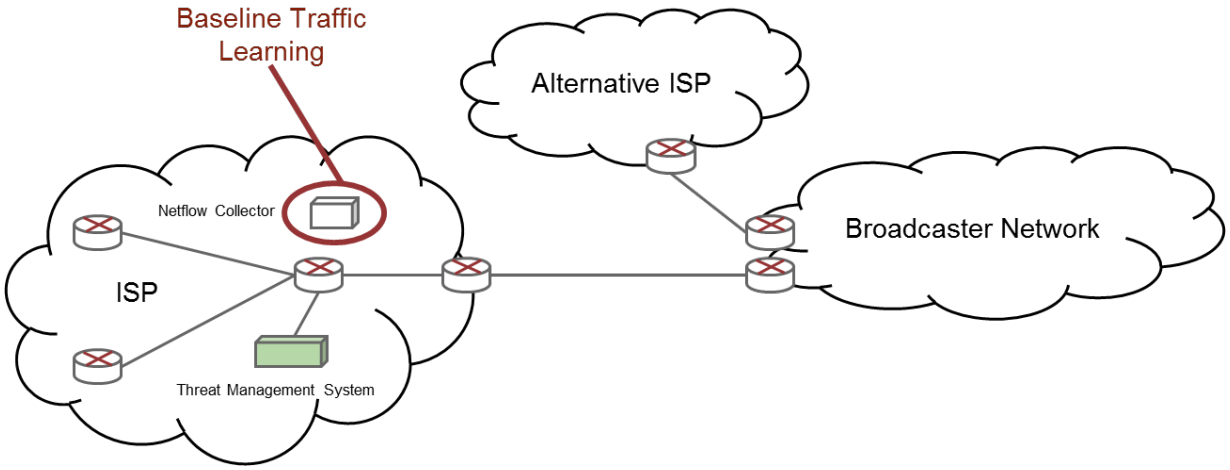

Figure 1: External ISP based solutions

For a broadcast network the baseline traffic may be too unspecific as a large download of a file based movie could be mis-detected as a DoS, or in case of a live sports event, streaming might consume much more traffic than that indicated by the baseline traffic profile.

It is best practice therefore to monitor endpoint devices and systems. These systems can be routers, firewalls but also http servers. If CPU consumption or routing tables look suspicious it might be a DoS or DDoS attack as the root cause.

For the incident and problem management process the Netflow Collector is just an additional source for information to aid in solving the problem. If the investigation indicates the presence of a DDoS attack the incident manager can access the TMS and activate specific mitigations to filter the traffic.

In the case of a dual homed setup (more than one ISP is used) the complete inbound and outbound traffic has to be rerouted through the network of the ISP that hosts the TMS system to mitigate the attack.

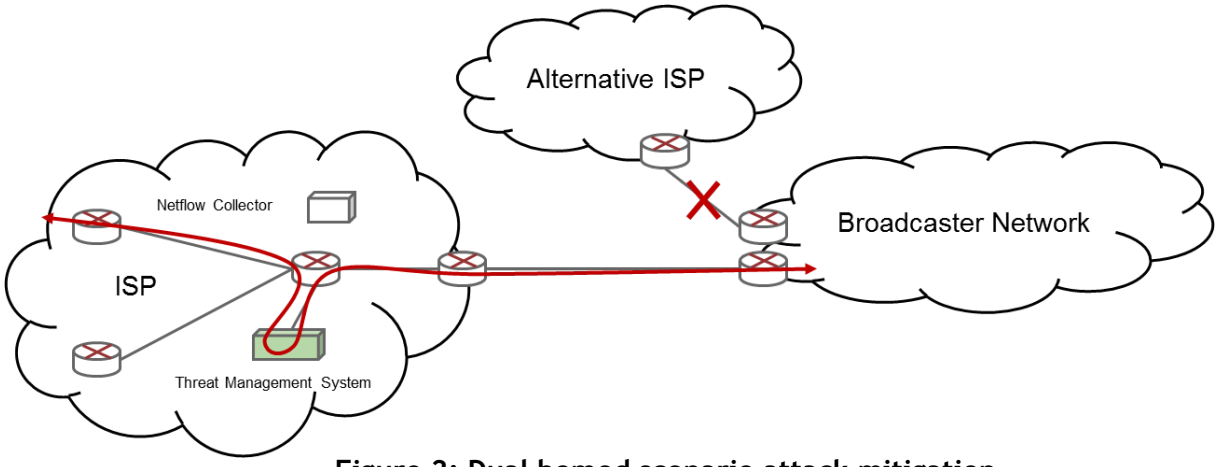

Figure 2: Dual homed scenario attack mitigation

Mitigation templates can be pre-defined to increase the effectiveness of filtering for the attack scenario. If mitigation is active, the broadcaster's network receives and sends clean traffic.

Before activating such a mitigation it has to be taken into account that good traffic might also be filtered and that the TMS system itself has a bandwidth limitation. The overall solution is better considered a "service continuity" solution as it keeps both the network and services that are inside the broadcaster's network online. It is not possible to permanently activate the TMS.

Pros:

The solution is externalized so the malicious traffic can be kept out and filtered before it trespasses the internal network perimeter. Additionally the external partner might have more knowledge in DDoS protection as the service is usually offered to several customers. This helps in reducing operational/personnel costs.

Cons:

As media traffic always varies, ISP-based solutions often produce a high proportion of false positives. The monitoring needs to be done on the service itself. Depending on the solution it can't be activated all the time; it is just activated to keep services alive in case of an incident. If there is a long-term DDoS hitting the broadcaster's network, additional countermeasures need to be developed and applied.

4.2.2 Solution 2 - On-premises solution

This solution is based on an on-premises DDoS mitigation appliance. The entire process of detection and 'washing' of traffic is done within the broadcaster's network. The setup consists of one appliance that is placed out of band at the Internet Edge layer of the network:

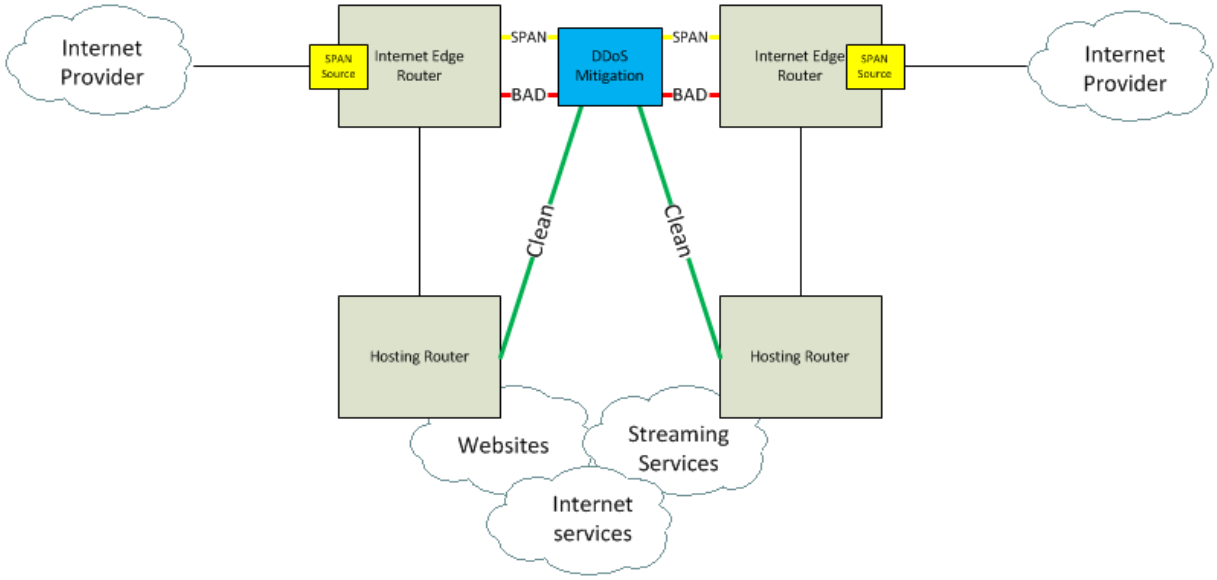

Figure 3: On-premises solution with internet Edge routers

Detection:

The appliance receives traffic information through SPAN ports. All the incoming traffic is mirrored to the Detection engine of the DDoS mitigation appliance. There is no need to mirror the outbound traffic because the incoming packets contain enough information to determine if there is a DDoS attack in progress.

The appliance keeps a profile of every IP address it protects. Using information such as 'maximum new connections per second' or 'maximum HTTP GETs per second', it creates a baseline for a specific service. As long as the traffic is below the baseline, there is no attack in progress. The appliance also checks for specific traffic patterns to check for 'low bandwidth attacks'.

Normal operation:

Because the traffic is only mirrored to the appliance, all the traffic still follows along its normal path. This means that the mitigation appliance is not critical to the day-to-day operations of the network, and can easily be serviced.

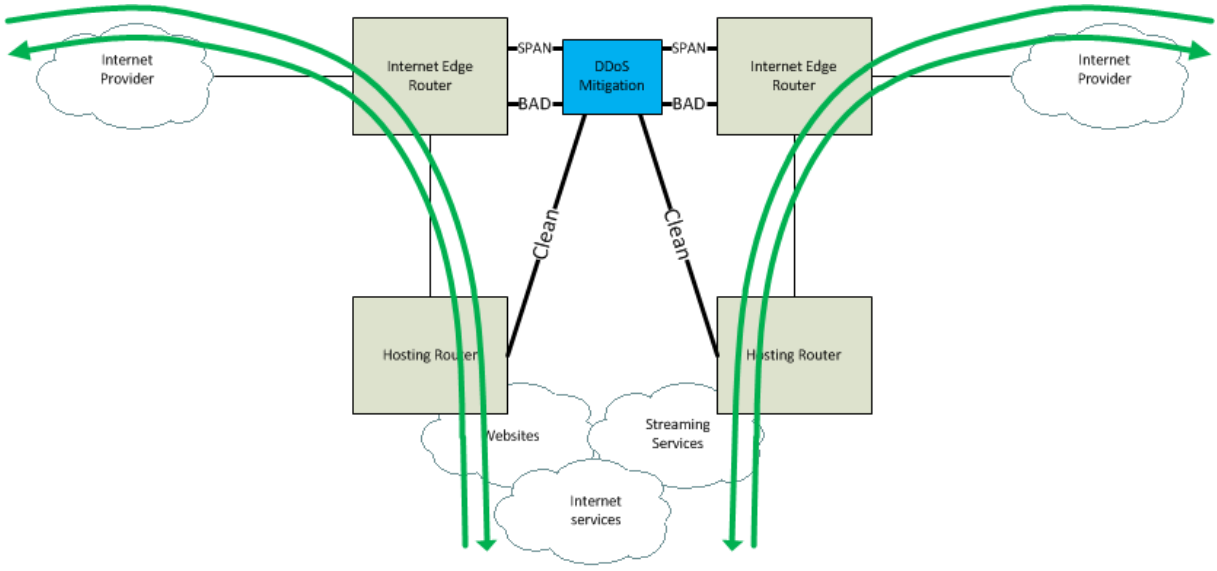

Figure 4: Normal traffic routing with appliance monitoring mirrored data

DDoS:

If the appliance detects a (possible) DDoS attack, it determines the IP address under attack, and sends a route update for that specific IP address to the Internet Edge routers through the BGP (Border Gateway Protocol). All the traffic for that IP address (both DDoS traffic and legitimate traffic) is routed from the Internet Edge router to the scrubbing engine of the appliance. The DDoS packets are discarded, and the clean traffic is routed to the destination service through separate, dedicated links.

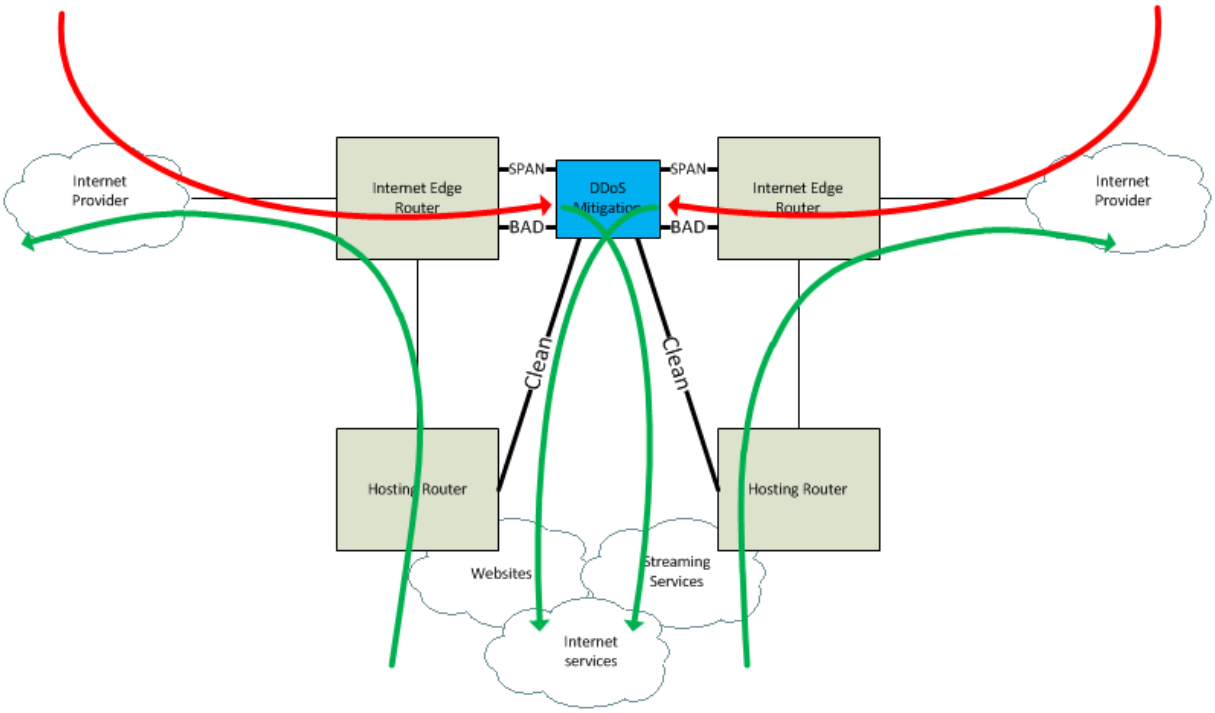

Figure 5: traffic routing through appliance during DDoS attack

If the appliance determines that the IP address is no longer under attack (certain thresholds are no longer exceeded), it withdraws the route update, and the traffic resumes its normal path.

Pros:

- In this setup, engineers have total control over the settings, thresholds and paths through the network.
- Because it is all internal traffic, engineers can specifically route 1 IP address through the appliance. External BGP-based solutions have to route an entire subnet of 255 addresses or more to receive the traffic for 1 IP address. This is inherent to the way Internet prefixes are exchanged from ISP to ISP.
- Out of band placement makes it easy to service the appliance.
- 'False positive' DDoS attacks (the appliance detects an attack, but it is actually all legitimate traffic) are being routed through the scrubbing engine, and to destination services without a big performance penalty because it is all internal traffic.

Cons:

- The potential DDoS is still being routed to the Internet Edge layer of the broadcaster. The available bandwidth at that layer must be more than the largest potential DDoS.

A hybrid setup, for example using this on-premises setup for day-to-day attacks and an external ISP based or third party solution for larger attacks can help in mitigating larger attacks.

4.2.3 Solution 3 – External third party based solution

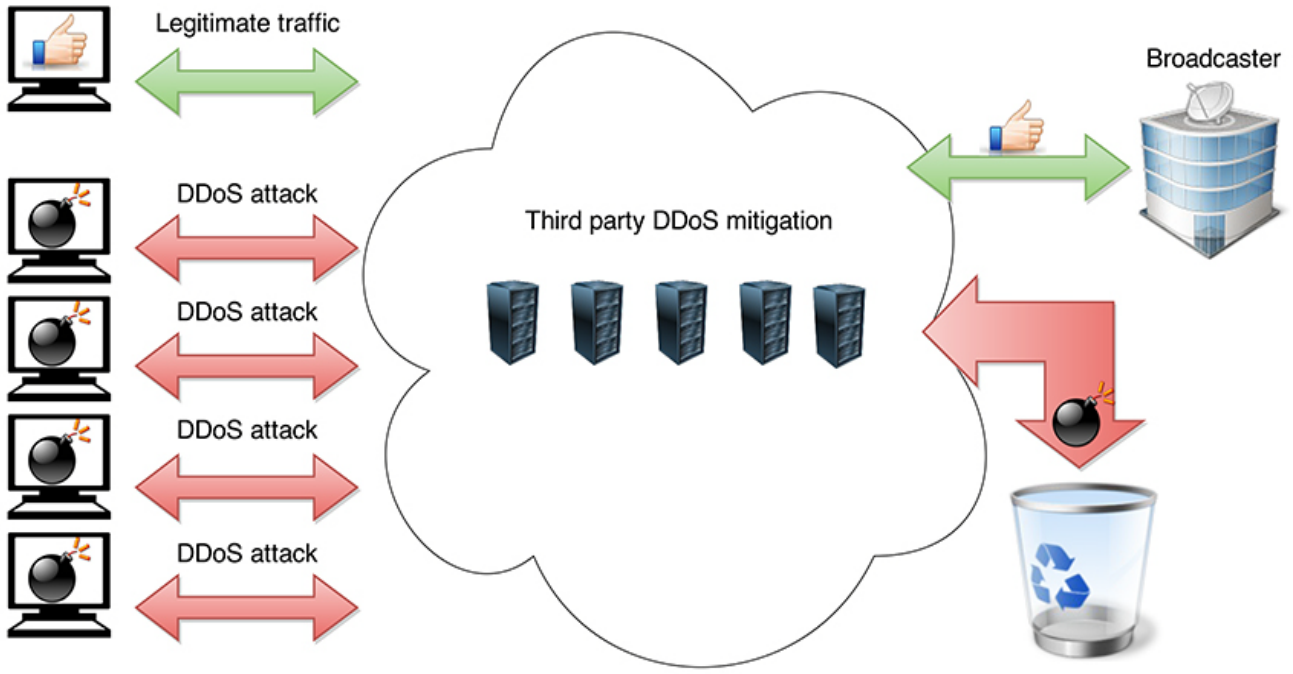

Figure 6: Third party mitigation

Another solution could involve renting the DDoS mitigation service from a third party provider with cloud based scrubbing centres in their backbone. DDoS attacks try to make a service or a network resource unavailable by leveraging coordinated actions from botnets of compromised computers. Such a provider can mitigate large amounts of DDoS floods in their own core network and hence ensure that valid traffic still flows as it should. Broadcasters are vulnerable when attacks of malicious traffic take up bandwidth, as it affects their ability to publish.

The vendor's Security Operations Centre detects network traffic anomalies by looking for typical signatures when monitoring traffic. DDoS attacks can be mitigated at every network layer, or typically ISO layers 3 to 7, depending on the vendor.

The customer can define which IP-ranges or subnets should be monitored. It is also important to define the levels of traffic where alarms will go off. The agreement must clearly state when the broadcaster shall be notified, and which contacts are to be alerted. It can also be agreed that mitigation starts automatically at certain thresholds.

Pros:

This solution with 'payment per mitigation' can be cost-effective compared to purchasing or renting equipment. Some attacks are so puny that there is no point in dealing with them by this means. Then again, some volumetric DDoS attacks could be too large to be mitigated on-premises.

Cons:

A disadvantage of renting the service is that it is a shared resource and not for the exclusive use of one customer. Having a separate scrubber locally or internally gives more control. Mitigation might also start faster if the broadcaster does the detection itself.