F5 White Paper

Intelligent Layer 7 DoS and Brute Force Protection for Web Applications

Both Denial-of-Service (DoS) and Brute Force Attacks have existed for many years, and many network devices tout the ability to withstand them. However, most of today's DoS attacks target layer 7 (L7) by overwhelming applications with seemingly valid requests and Brute Force programs can send more than one million attempts per second. This paper will discuss how to intelligently mitigate these types of attacks.

by Michael Koyfman
Sr. Field Systems Engineer

Peter Silva
Technical Marketing Manager

Contents

Introduction

Denial-of-service (DoS) attacks have existed for many years. A variety of motives drive these attacks—from script kiddies declining service to a target website to a well-organized crime syndicate using DoS attacks for political warfare or monetary extortion. There are also unintentional attacks, such as increased website popularity or simply the "Slashdot effect" (where a high traffic site links to a smaller site not equipped to handle the increased load). Attacks can also result from a search engine indexing the website during peak times.

Because of the variety of motivations and desired outcomes, you must take various factors into account when detecting a true layer 7 DoS attack. Traditionally, most DoS attacks are carried out at layer 3 and 4, overwhelming a particular IP address and/or port with an inordinate amount of either legal traffic or malformed requests that exhaust available resources (the number of available ephemeral ports, the number of concurrent layer 4 connections, available bandwidth, and more). Many network devices—such as firewalls, routers, and load balancers—offer a certain degree of protection against those types of DoS attacks, since layer 3 and 4 malicious behavior is fairly easy to identify and block. However, properly detecting layer 7 DoS attacks requires intelligence, historical analysis, and inspection of each request. Without an effective layer 7 attack detection and prevention mechanism, websites and applications are unable to offer high availability and overload protection.

Detecting a Layer 7 DoS Attack

There are several challenges that arise when applying traditional solutions to layer 7 DoS attack mitigation. These challenges can range from detecting and stopping the attacker's requests without user impact to accurately triggering the beginning of the prevention stage. Traditional solutions, such as Intrusion Prevention Systems (IPS), often limit the number of requests from a source IP address during a designated period of time. While this certainly helps mitigate attacks, it can unreasonably restrict the server's ability to process valid requests. A single source IP can also represent multiple clients, creating a high degree of false positives from source IP address-only mitigation.

In order to intelligently and accurately detect a layer 7 DoS attack, you need to inspect either the server response time or the request frequency rate for abnormalities, and take a close look at latency. Under the normal load, the

response time for each request should be similar. While some variance certainly exists, any significant increase in the application response latency could indicate an unbearable load. As such, increased latency should be used as a trigger step for detecting a layer 7 DoS attack.

Because layer 7 DoS attacks can target a particular resource—like a CPU-intensive, business-logic code of an application (such as re-sizing images or requesting a large file)—the latency triggers need to be enforced on per-object level to block requests to the affected objects while allowing traffic to objects that do not experience response time degradation. For example, a 100 percent increase in the application response time might be considered abnormal, and additional examinations would need to be performed to mitigate the situation. Applications might also have a minimum acceptable latency regardless of heavy load, so a minimum latency threshold should also be taken into account when investigating the suspicious behavior further.

Following the detection of a significant latency increase, it is important to determine whether you need further action. After examining the increase in the requests/sec on the IP-based level or object-based level by comparing these numbers with past activity, you can identify suspicious versus normal latency increases.

Mitigating Layer 7 DoS Attack

Once the suspicious behavior has been analyzed and identified as a potential attack, some preventative measures can be initiated. There are several approaches to successfully stopping attackers in their tracks. One method is to inject a small piece of JavaScript in the server response. That code needs to be evaluated by the browser and then the code assigns a specific value to the application firewall. This method verifies that the request is valid, coming from a real browser instead of a script. This method accurately distinguishes malicious requests from the legitimate ones. Because many of the attacks are script-based—since browsers were not designed efficiently to send several requests/second—this method provides a high degree of efficiency and accuracy in defending against layer 7 DoS attacks. While this can alleviate potential botnets, the back-end application might still be at risk. To avoid application overload, prevention methods can also include throttling requests/sec to a certain object or to number and limiting a specific number of requests/sec based on source IP address. These can also be used if deemed appropriate, based on the situation and configuration. These methods can also be combined to provide the most efficient attack mitigation on a single device.

F5® BIG-IP® Application Security Manager™ (ASM) version 10 (v10) is equipped with a unique technology that provides a new and effective approach to protecting web applications. Specific to DoS prevention, BIG-IP ASM v10 can configure detection criteria for server latency and transactions-per-second (TPS.) If any attacks trigger your set latency thresholds, BIG-IP ASM marks the IP address as suspicious. At that point, the DoS engine starts collecting TPS information for the suspicious IP and analyzes TPS history rate to determine if it is a real DoS attack or a false positive. If the suspicious criteria reach its TPS, the engine will start dropping connections by either source IP or per URL, depending on the policy. When the engine in suspicious criteria mode reaches the TPS (per IP address), the prevention policy moves to Source IP-Based Rate Limiting mode. If the TPS (per URL) is reached, BIG-IP ASM will apply URL-Based Rate Limiting. BIG-IP ASM will then reset those connections. If the suspicious criteria still considers the request as an attack, BIG-IP ASM will try the second and third policy with a two minute interval between them.

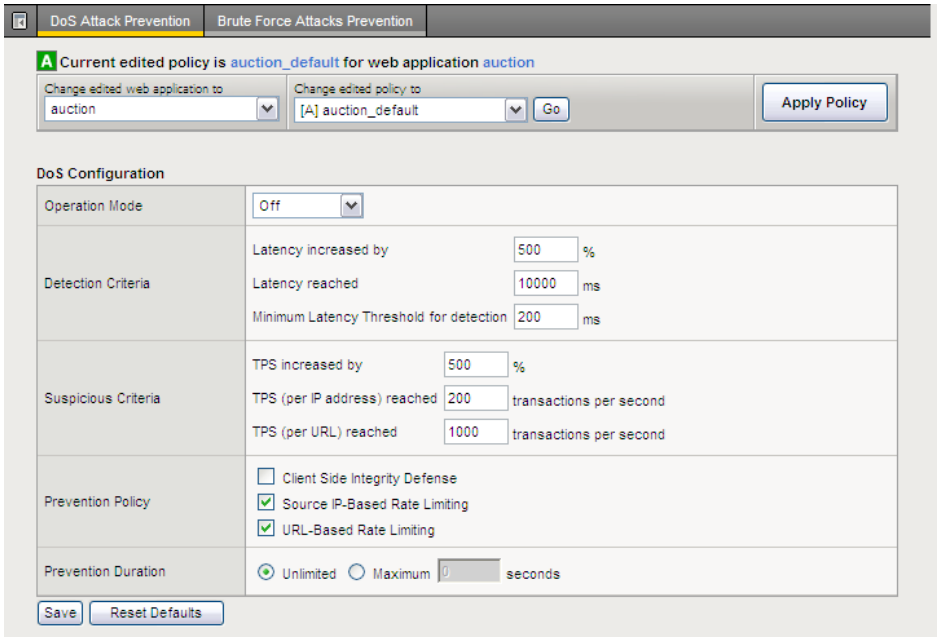

BIG-IP ASM administrative GUI showing DoS configuration.

You can also set a Prevention Duration. This limits the connection, dropping it to certain number of seconds or leaving it as long as the attack occurs, with no time limit. Finally, the reporting page for the DoS engine provides the values it detected so you can make the necessary adjustments to prevent false positives while learning more about any other DoS attacks. The report also shows you which standard prevention policy was applied and which one was previously applied and didn't mitigate the attack.

Protecting against Brute Force Attacks

A similar method can be applied when protecting one of most common and targeted area of the application: the login and authentication page. Very often, a DoS attack comes in the form of a Brute Force Attack on the application login page and attempts to "guess" the login credentials, gaining unauthorized access to the application. The same principles for layer 7 DoS attack prevention can be applied to the application's login page.

Many applications already disable login attempts based on the number of failed attempts per user account, but they usually lack visibility into general Brute Force Attack behavior that might be directed at it, such as guessing passwords for many different user accounts sequentially. Throttling the maximum number of failed login attempts on a per browser-session and per-IP address level, along with taking abnormal failed login increases into account, can protect an application against Brute Force Attacks while enabling legitimate users to log in.

BIG-IP ASM v10 gives you the ability to define a URL (or set of URLs) that will enforce the Brute Force and access validation criteria for a successful or failed login. BIG-IP ASM will track the defined URL and decide when a Brute Force Attack is occurring. BIG-IP ASM adds any listed URLs to the security policy and will then search for the URL(s) or the parameter names in each request to take a measurement on the request rate.

Similar to the DoS protection engine, you can set detection criteria on BIG-IP ASM; these are measured by a percentage of failed login rates or by failed logins per second. Once threshold for detection criteria are reached, the IP address is considered suspicious and the engine will collect failed login attempts per this specific IP. If it reaches the suspicious criteria, the protection engine will then apply the prevention policy.

You can also define the criteria for a successful login or a failed login by time period. For instance, if the client fails to login five times (or had five attempts within a specified time frame), they will have to wait 10 minutes until they can login again. You can also set the protection engine to prevent any future logins from a specific IP address. With BIG-IP ASM, you have various detection options that can be arranged as one string or multiple triggers; all it takes is one valid statement to activate the prevention policy. This means that if you have a few

options configured and only one of them is triggered, BIG-IP ASM will rely on these criteria for decision making.

The Brute Force engine also has a reporting page to indicate which thresholds were triggered, how many connections were dropped, when the prevention was executed, and how long the brute force attempt occurred. The report also provides information on which prevention policy was applied during the attack along with the source IP from which it came.

The DoS and brute force prevention engines are powerful tools to deal with the ongoing threat of layer 7 attacks, but it is important to verify any false positives before switching to blocking mode in production environment.

Conclusion

Layer 7 DoS attacks are a real threat to all Internet-exposed sites. They threaten to disrupt the functionality and services that these sites and applications provide, as well as allow unauthorized access to the protected information. Their detection and mitigation is a non-trivial task that requires a careful, intelligent approach to effectively detect and block only the malicious requests without disrupting legitimate users. An application firewall, such as F5 BIG-IP ASM v10, is perfect place to detect such attacks and protect the application and no other web application security solution offers the detailed control and protection from both layer 7 DoS and Brute Force Attacks.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

CS97261-WP-ASM v10 04/09