

Slow Client Attack Prevention

<https://campus.barracuda.com/doc/3569/>

The Slow Client Attack Prevention feature is available ONLY in Firmware Versions 7.7 and above.

Overview

In a slow client attack, an attacker deliberately sends multiple partial HTTP requests to the server to carry out an HTTP DoS attack on the server. The client attempts to slow the request or response so much that it holds connections and memory resources open on the server for a long time, but without triggering session time-outs. Common ways to carry out this attack include:

- **Slow HTTP Headers Vulnerability (Slowloris)** – As described in Slowloris HTTP DoS (<http://ha.ckers.org/slowloris/>), using this technique the client never completes sending the headers. It sends headers one-by-one at regular intervals to keep sockets from closing and the web servers thereby tied up. In particular, threading servers tend to be vulnerable when they try to limit the amount of allowed threading. Slowloris must wait for all of the sockets to become available before successfully consuming them, so for high traffic websites, it may take awhile for the site to free up its sockets.
- **Slow HTTP POST Vulnerability (R-U-Dead-Yet or RUDY)** – Using this technique, the client attempts to DoS the server using long form field submissions. The client sends all of the HTTP headers, one of which is a legitimate Content-Length header with a large value. The client then iteratively injects data into the form's post field at a very slow rate, so the web application keeps waiting for the full data to arrive. Once multiple threads are tied up by waiting, the server eventually runs out of resources and gets DoS'ed. More technical details about layer-7 DDoS attacks can be found in the OWASP lecture: OWASP-Universal-HTTP-DoS (<http://www.hybridsec.com/papers/OWASP-Universal-HTTP-DoS.ppt>).
- **Slow Read DoS Attack** – Using this attack technique, the client request completes fully. When the server responds, the client advertises very small windows for accepting response data. For a large response (a file download, for example) the client's slow reception rate ties up server resources for a long time. Multiple requests of this type can eventually take the server down.

These requests are layer 7 DoS attacks. They are typically legitimate from a protocol compliance point of view and are therefore not detected by network layer DDoS devices, by IPS/IDS, or even by your ISP. Clients can DoS the server stealthily and slowly, without consuming any significant bandwidth on the network, so they remain otherwise undetected.

The **WEBSITES > DDoS Prevention** page allows you to configure slow client attack prevention for HTTP and HTTPS Services.

How does Slow Client Attack Prevention Work?

The following settings allow the identification of prevention of a slow client request or response attack:

Data Transfer Rate

The minimum data transfer rate the Barracuda Web Application Firewall expects for requests from the client and responses to the client. Data transfer rates slower than this are considered slow.

Max Request Timeout

The maximum time allowed to receive a request from a client. If a request does not complete in this time, the connection is terminated, FIN is sent to the client, and further requests are blocked.

Max Response Timeout

The maximum time allowed to send a response to the client. If the response transfer is not complete in this time, the connection is terminated, Fin is sent to the client, and further responses to the client are not sent.

Incremental Request Timeout

This value specifies the initial timeout window a client has in which to complete a request. The system then progressively shrinks the window using an adaptive algorithm. If the client repeatedly fails to complete a request in the shrinking window, the request timeout window converges to zero and the connection is dropped. If the client begins to send data at a healthy rate, the window is progressively expanded.

This adaptive algorithm ensures that temporary network delays do not affect genuine clients, but persistent slow clients are detected and denied.

Incremental Response Timeout

This value specifies the initial timeout window a client has in which to receive a response. The system then progressively shrinks the window using an adaptive algorithm. If the client repeatedly fails to receive the response in the shrinking window, the response timeout window converges to zero and the connection is dropped. If the client begins to receive data at a healthy rate, the window is progressively expanded.

This adaptive algorithm ensures that temporary network delays do not affect genuine clients, but persistent slow clients are detected and denied.

Exception Clients

The IP addresses that should be exempted from slow client attack prevention. Specify a single IP address or range of IP addresses, or a combination of both using a comma delimiter with no spaces.

Steps to Configure Slow Client Attack Prevention

To view or edit Slow Client Attack Prevention for a Service, perform the following steps:

1. From the **WEBSITES > DDoS Prevention > Slow Client Attack Prevention** section **Edit** the Service requiring the protection.
2. In the **Edit Slow Client Attack Prevention** page, you can view or edit the configured values.
3. Click **Save** after modifying values. For more information, click **Help** on the web interface.

© Barracuda Networks Inc., 2025 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.