



How to Configure Rate Control

Rate Control

The Rate Control policy defines shareable policies for controlling the rate of requests to a web application. The Client/HTTP requests are delivered to an application server, or parts of it, which can get overloaded with peak traffic and create high response at times. The Barracuda Web Application Firewall's Rate Control policy prevents application servers from being overloaded. A Rate Control Pool specifies the maximum number of Active Requests and Client Backlogs along with a set of Preferred Clients. A Preferred Client specification defines a range of IP addresses and an associated weight. The Barracuda Web Application Firewall uses these weights to perform a weighted round robin scheduling between queues when forwarding requests to the application server from the rate control pool. You can set a Rate Control pool to limit the client requests. The Barracuda Web Application Firewall collects requests and queues them in the pool. The back-end application servers receive requests from the pool at the rate specified.

A *default* rate control pool is provided by the Barracuda Web Application Firewall which allows the easy set up of Rate control. Using **Edit** on the **BASIC > Services** page or using **Edit** on the **WEBSITES > Advanced Security** page, you can select a rate control pool to apply to your Service or URL policy. **Edit** the desired Service or URL Policy to select a Rate Control Pool. A custom Rate Control Pool can be created on the **ADVANCED > Libraries > Rate Control Pool** page by selecting **Add Rate Control Pool**.

Before you set up a Rate Control Pool

Answer the following questions:

1. What are the maximum simultaneous requests that can be served by the resource being protected? This determines the **Maximum Active Requests** setting.
2. What, if any, are the bonafide gateways and mega-proxies that will be accessing the protected resources? These are **Preferred Clients**. If they proxy client requests, assign a suitable weight to the proxy IP address and if they relay a set of client IP addresses, then assign a weight to the range of IP addresses.
3. What is the maximum queue to allow for IP addresses not defined in Step 2? This defines the **Maximum per client backlog** setting.

Benefits of a Rate Control Pool

A Rate Control Pool helps defend against rate control attacks by:

- Throttling attackers attempting to flood the application with DoS attacks. The requests get queued for weighted round robin scheduling, slowing down the request rate seen by the server.
- Protecting "Load Sensitive" applications, such as search or DBMS intensive applications, from application DoS attacks.
- Allowing bonafide gateways and mega-proxies access while preventing attacks.

Scheduling algorithm for Rate Control Pool

The scheduling algorithm between queues is weighted round robin. Implicitly, the weight of each unconfigured client queue is 1. For example, a Preferred client is defined with weight 5 and at a given time the Barracuda Web Application Firewall has queues for 2 Unconfigured clients with a few requests in each. The Barracuda Web Application Firewall will serve 1 request from each unconfigured client queue followed by 5 requests from the Preferred client queue.

Rate control policies can be specified per service or per URL policy. Rate Control Pools are defined on the **ADVANCED > Libraries** page. These rate control pools are globally shareable among services or among URL



policies or both. Once defined, they can be bound to multiple services on the **BASIC > Services** page, when you **Edit** a service. Also they can be bound to multiple URL policies on the **WEBSITES > Advanced Security** page, when you **Edit** a URL policy.

Customizing Rate Control Pool

From the **ADVANCED > Libraries > Rate Control Pool** page use **Edit** or **Add Rate Control Pool** to customize a rate control pool. Set the following values:

- **Rate Control Pool Name** – Enter a name for the new rate control pool.
- **Maximum Active Requests** – Enter the maximum number of Active Requests processed at a given time by the Barracuda Web Application Firewall. An active request is a request which has not fully completed.
- **Maximum per client backlog** – Enter the number of requests per client IP address that will be queued when the Barracuda Web Application Firewall has reached the **Maximum Active Requests** limit. For example, if **Maximum Per Client Backlog** is set to 32 and the Barracuda Web Application Firewall is processing the default 100 **Maximum Active Requests**, then for any given client IP, the Barracuda Web Application Firewall will queue up to 32 requests. Any requests after that will be dropped until a request is deleted from the queue.
- **Maximum Unconfigured Clients** – Enter the maximum number of Unconfigured Clients. All clients which are not Preferred Clients are Unconfigured Clients. For each unique client IP, the Barracuda Web Application Firewall will maintain an individual backlog queue. For example, if **Maximum Unconfigured Clients** is set to 100 and **Maximum Per Client Backlog** is set to 32, the Barracuda Web Application Firewall will maintain 100 queues each with 32 pending requests, a total of 3200 pending requests.

Click **Add** to save the configuration.

Use **Add Preferred Clients** to add a single or range of IP addresses to the pool gets preferred treatment. Each Preferred client has a configured **Weight** and its own queue containing **Max Per Client Backlog** times **Weight**. For example, if **Max Per Client Backlog** is set to 32 and preferred client **Weight** is set to 5, then the queue size will be 32 x 5. If the preferred client queue contains a range of IP addresses, the queue will include all requests from all the clients falling within that range.

Creating Preferred Clients

Click **Add Preferred Clients**, under **Options**. The Add Preferred Client window appears. Specify values for the following fields:

- **Name** – Enter the name for the client weight.
- **Status** – Sets the status of the preference. Enabling this makes the client IP address range a preferred list of IP addresses.
- **Preferred Client IP Range** – Enter the IP address or the range of IP addresses (For example: 10.0.0.1 – 10.0.0.10) which will be treated in a preferential manner. Preferred Client is an IP address or a range of IP addresses with an associated weight.
- **Weight** – Enter the weight for the range of IP addresses. These IP addresses are evaluated in the order of their weights; the higher the weight the higher the precedence (1 is the lowest priority and 100 the highest priority).

Click **Add** to save the configuration.

Click **Delete** to delete the created Rate Control Pool from the list.

