
Global Server Load Balancing Overview

<https://campus.barracuda.com/doc/19333829/>

Overview

Global Server Load Balancing (GSLB) allows you to coordinate how traffic is processed among multiple data centers. A Barracuda Load Balancer ADC acts as a controller, selecting the location to which traffic is directed based on the parameters that you configure and the health of the data centers. This allows you to allocate the work among multiple data centers and to ensure that if one data center fails then traffic is redirected automatically to a functioning data center.

GSLB Examples

GSLB can be useful when you have:

- a number of server farms that are physically located around the world and you want incoming connections to be directed to the closest healthy server farm.
- two data centers and you want one of them to be reserved for use in the event of a disaster. You can assign the first with a high priority and have all traffic directed to it, while the other is used only if the first data center fails.
- multiple data centers and each has region-specific content. Depending on the location of the client, requests can be directed to the data center most appropriate for that region.

GSLB Definitions

- **Site** – a network location that hosts data. It may be a Service on a Barracuda Load Balancer ADC with a server farm or one Real Server.
- **GSLB Controller** – the Barracuda Load Balancer ADC which determines where traffic is directed. It contains configuration information about the sites and it performs health checks on all sites in regular intervals. Only one GSLB Controller is active at a time. It is recommended that you configure one or more backup GSLB Controllers.
- **Region** – defines a geographical area, usually composed of one or more countries. You can define custom regions or use the predefined regions.

How GSLB Works

Figure 1 shows how the GSLB devices at each site are configured as authoritative nameservers for the

domain www.mydomain.com. They also perform local load-balancing duties for their own four web servers.

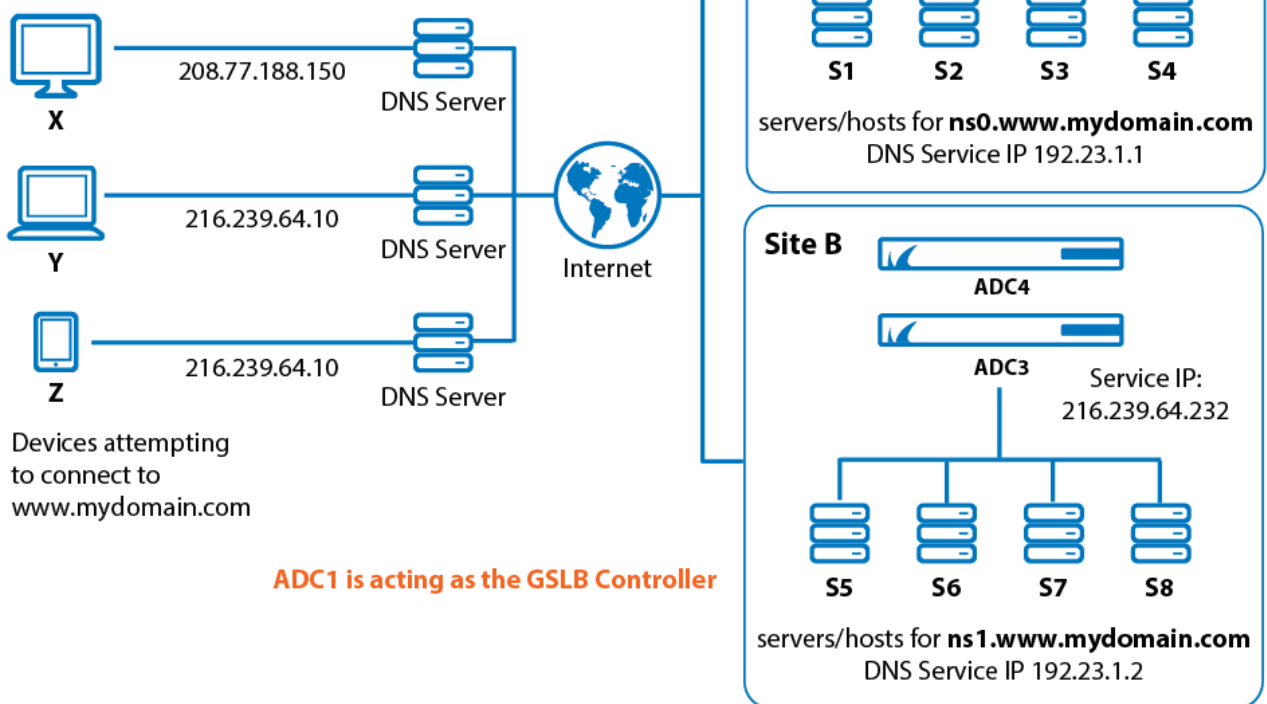
As long as both sites and services are healthy, when a client makes a request for www.mydomain.com they are directed to the clustered GSLBs in Site A. The GSLBs in Site A respond with an IP list based on the selected algorithm, described below.

Figure 1. GSLB devices configured as authoritative nameservers and load balancers.

DNS Server configuration might be:

```
www.mydomain.com IN NS ns0.www.mydomain.com
                      ns1.www.mydomain.com
```

```
ns0.www.mydomain.com IN A 192.23.1.1
ns1.www.mydomain.com IN A 192.23.1.2
```



The following figure shows the DNS Service IP Address for ADC 1, specified on the **TRAFFIC > GSLB Settings** page.

DNS Service

| | |
|------------------------|---|
| Name | <input type="text" value="dns"/> |
| DNS Service IP Address | <input type="text" value="192 . 23 . 1 . 1"/> <small>IP address on which the DNS will respond.</small> |
| Netmask | <input type="text" value="255 . 0 . 0 . 0"/> |
| Interface | <input type="text" value="ge-1-1"/> |

Site Selection Criteria

The GSLB Service allows you to specify traffic to be directed to a site based on one of three parameters:

- **Geo IP** – Proximity of the system making the request to a site that can serve the request;
- **Region Only** – Region of the system making the request; or
- **Priority** – Priority order of the sites.

To provide location-based Response Policies, the Barracuda Load Balancer ADC uses a database of IP addresses and geographical locations. This database is updated by the Location Definitions, which are part of the Energize Updates maintained by Barracuda Central.

The following sections consider each of these site selection criteria, using the diagram above.

Geo IP

The GSLB Controller determines the location of the system making the request based on the Location Definitions and compares that to the location of each site. It returns a list of site IP addresses ordered from closest to furthest.

You might choose to use Geo IP if you have a number of server farms that are physically located around the world, and you want clients to be directed to the closest healthy server farm. Set the **Response Policy** to **Geo IP** to send client requests to the geographically nearest site. If you have a backup site, set the **Failover IP Address** to its IP address.

Geo IP does not consider site priority.

In the diagram above, Device Z is closest to Site B. So traffic from Device Z would be routed through

the internet to Site B.

Following is a portion of the **TRAFFIC > GSLB Services** page, showing configuration for **Geo IP**.

| | | | | | | | | | |
|--------------|---|-----|--------|---|----------------|--|------|-------------|---------|
| mydomain.com | | | | | | | Host | Delete | |
| www | A | www | Geo IP | | | | Site | Edit Delete | |
| | | | | ✓ | 208.77.188.166 | | | Edit Delete | Disable |
| | | | | ✓ | 216.239.64.232 | | | Edit Delete | Disable |

Region Only

The GSLB Controller determines the region of the system making the request based on the Location Definitions.

- If the originating system is in a region that is associated with one or more sites, a list of the healthy site IP address(es) is returned. The most specific matches appear first in the list; any sites associated with All Countries are last in the list.
- If the location of the originating system cannot be determined, any healthy sites that are associated with All Countries are returned.
- If neither of the preceding cases identifies at least one site IP address, the Failover IP Address is returned.

You might choose to use Region Only if you have multiple data centers, each with region-specific content, and you want client requests from a certain region to be directed to the data center that supports that region. Set the **Response Policy** to **Region Only**, to associate requests with a region based on the location of the client, and direct traffic to the appropriate data center. If you have a backup site, set the **Failover IP Address** to its IP address. You can use content switching rules to direct HTTP traffic within the backup data center.

Region Only does not consider site priority.

In the diagram above, Device X is in the region with Site A. Devices Y and Z are in the region associated with Site B. Traffic from Device X will go to Site A. Traffic from Devices Y and Z will go to Site B.

Following is a portion of the **TRAFFIC > GSLB Services** page, showing configuration for **Region**.

| | | | | | | | | | |
|--------------|---|-----|-------------|---|----------------|------------------------|------|-------------|---------|
| mydomain.com | | | | | | | Host | Delete | |
| www | A | www | Region Only | | | | Site | Edit Delete | |
| | | | | ✓ | 208.77.188.166 | Regions: All Countries | | Edit Delete | Disable |
| | | | | ✓ | 216.239.64.232 | Regions: Canada | | Edit Delete | Disable |

With the configuration shown here, when the client from Canada makes a DNS request to www.mydomain.com, the GSLB controller returns the Canada region IP 216.239.64.232. If the client

does not fall in Canada region, it will return the default All Countries IP 208.77.188.166.

Priority

The GSLB Controller returns a list of site IP addresses ordered from lowest to highest priority value.

Priority does not consider location.

Priority is often used for disaster recovery. As in the diagram above, you have two sites and you want all traffic directed to Site A while Site B is on standby and used only in the case of the failure of Site A. Create an entry for each site giving the Site A **Priority 1** (highest) and Site B **Priority 2**. Set the **Response Policy** to **Priority** so only priority is considered when directing traffic.

When a query for the address of the domain name is received, a response containing one or more IP addresses is returned. If it is operational, the primary site's (Site A's) IP address will be returned first in the list and the backup site's (Site B's) IP address will be second. If the primary site becomes unavailable, only the second site's IP address will be returned.

The primary site will be monitored, even after failure, so that when it becomes available, its IP address will once again be first in the returned list.

Following is a portion of the **TRAFFIC > GSLB Services** page, showing configuration for **Priority**.

| | | | | | | | | | | |
|--------------|--|---|-----|----------|---|----------------|-------------|------|-------------|---------|
| mydomain.com | | | | | | | | Host | Delete | |
| www | | A | www | Priority | | | | Site | Edit Delete | |
| | | | | | ✓ | 208.77.188.166 | Priority: 1 | | Edit Delete | Disable |
| | | | | | ✓ | 216.239.64.232 | Priority: 2 | | Edit Delete | Disable |

Defining Response Policy

Define the Response Policy on the **TRAFFIC > GSLB Services** page. Select Geo IP, Priority, or Region.

- If you select **Priority**, define the order of priority on the same page.
- If you select **Region** or **Geo IP** settings, navigate to the **TRAFFIC > GSLB Settings** page.
 - Geo IP** – Enter the IP address of the server, so the
 - Regions** – Several predefined regions are listed on the **TRAFFIC > GSLB Settings** page. You can also create a custom region on that same page. Scroll down to **GSLB Regions** and click **Add Custom Region**. Name the custom region and select one or more small regions from the list. Then click **Save**. Return to the **TRAFFIC > GSLB Services** page to specify using the new Custom Region.

Refer to [Installation of Global Server Load Balancing](#) for instructions on how to install multiple GSLB Controllers.

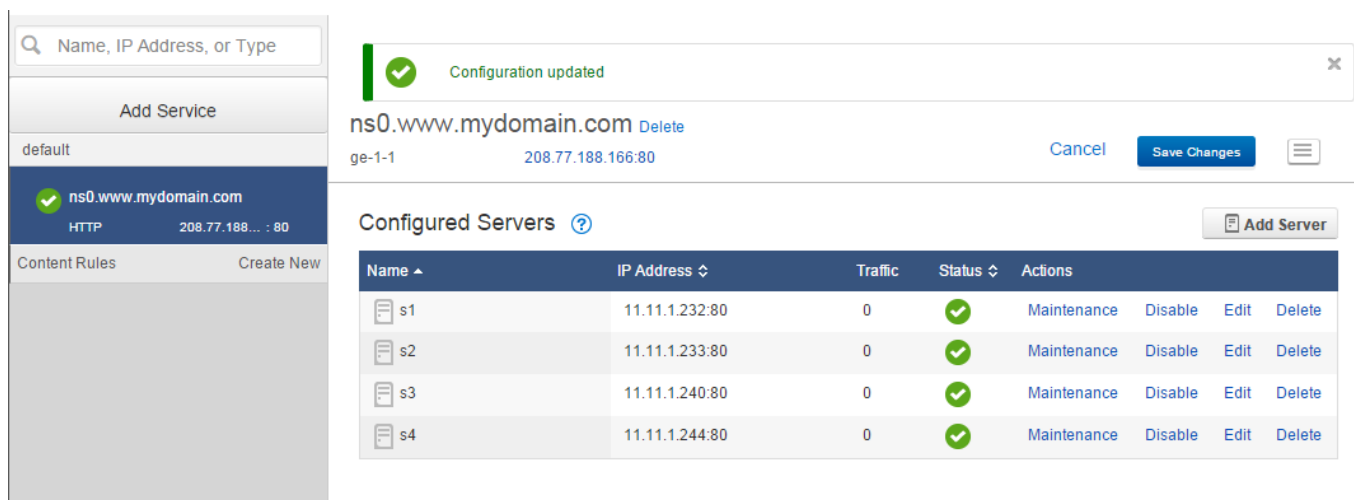
Configuring Multiple GSLB Controllers

Only one Global Server Load Balancing (GSLB) Controller is active at any one time. However, you can configure multiple GSLB Controllers to increase the availability of your infrastructure in these two ways:

- Operate in [High Availability](#) mode, in which case all of the GSLB information is copied to the passive system.
- Configure one or more other Barracuda Load Balancer ADCs (or clustered pairs) as GSLB Controllers where:
 - Each system or clustered pair has a DNS entry pointing to it. The first available entry is used by a client.
 - The GSLB configuration is synchronized manually between all GSLB Controllers unless they are passive systems in a cluster.

Configuration

This screen shows configuration for the Service IP on Site A for the domain [ns0. www.mydomain.com](#).



The screenshot displays the configuration interface for Site A. On the left, a sidebar contains a search bar, an 'Add Service' button, a list of services with 'ns0.www.mydomain.com' selected, and a 'Content Rules' section. The main area shows a confirmation message 'Configuration updated' and the selected service details: 'ns0.www.mydomain.com' with a 'Delete' link, 'ge-1-1' as the interface, and '208.77.188.166:80' as the IP address. Below this is a 'Configured Servers' table with an 'Add Server' button.

| Name | IP Address | Traffic | Status | Actions |
|------|----------------|---------|--------|---------------------------------|
| s1 | 11.11.1.232:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s2 | 11.11.1.233:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s3 | 11.11.1.240:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s4 | 11.11.1.244:80 | 0 | ✓ | Maintenance Disable Edit Delete |

This screen shows configuration for the Service IP on Site B for the domain [ns1. www.mydomain.com](#).

[Add Service](#)
default

ns1.www.mydomain.com

HTTP 216.239.64... : 80

[Content Rules](#) [Create New](#)

Configuration updated

ns1.www.mydomain.com [Delete](#)
ge-1-1 216.239.64.232:80 [Cancel](#) [Save Changes](#)

Configured Servers ?

[Add Server](#)

| Name | IP Address | Traffic | Status | Actions |
|------|-----------------|---------|--------|---|
| s5 | 10.11.30.232:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s6 | 10.11.30.233:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s7 | 10.11.30.240:80 | 0 | ✓ | Maintenance Disable Edit Delete |
| s8 | 10.11.30.244:80 | 0 | ✓ | Maintenance Disable Edit Delete |

Failover

If no sites match the Response Policy or if all sites that match the Response Policy fail the health check, a pre-configured Failover IP address for the sub-domain is returned. This is the IP address of a site that can accept the traffic if the other systems become unavailable.

The health of the site at the Failover IP address is not monitored.

Monitoring

Servers must be monitored for their health. For the Monitoring IP address, enter the IP address of the server to be checked. If you will just use the site IP address for health checks, then leave this field blank.

Figures

1. GSLBexample2.png.png
2. DNSservice.png
3. config2_geoip.png
4. config3_region.png
5. config4_priority.png
6. serviceIP_siteA.png
7. serviceIP_siteB.png

© Barracuda Networks Inc., 2024 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.