

BIG-IP[®] DNS: Load Balancing

Version 12.1



Table of Contents

About Global Server Load Balancing.....	5
Introducing BIG-IP DNS.....	5
About global server load balancing.....	5
Static load balancing methods.....	5
Dynamic load balancing methods.....	7
About load balancing and resource availability.....	9
About virtual server dependency.....	9
Configuring virtual server availability to be dependent on the status of other virtual servers.....	9
Limit settings for resource availability.....	10
About wide IP-level load balancing.....	10
About the Global Availability load balancing method.....	10
About the Ratio load balancing method.....	11
About the Round Robin load balancing method.....	12
About pool-level load balancing.....	12
About the Drop Packet load balancing method.....	13
About the Virtual Server Score load balancing method.....	13
About the Virtual Server Capacity load balancing method.....	13
About the Round Trip Times load balancing method.....	13
About the Packet Rate load balancing method.....	13
About the Least Connections load balancing method.....	13
About the Kilobyte/Second load balancing method.....	13
About the Hops load balancing method.....	14
About the Completion Rate load balancing method.....	14
About the CPU load balancing method.....	14
About the Return to DNS load balancing method.....	14
About Static Persist load balancing.....	14
About the Fallback IP load balancing method.....	15
About the None load balancing method.....	15
About the QoS load balancing method.....	15
About dynamic ratio load balancing.....	18
Using the preferred load balancing method when metrics are unavailable.....	18
Configuring the resources in a pool for manual resume.....	18
Restoring availability of a pool member manually.....	19
Best Practices: Dynamic load balancing and performance configuration.....	19
Using Topology Load Balancing to Distribute DNS Requests to Specific Resources.....	21
How do I configure BIG-IP DNS to load balance DNS requests to specific resources?.....	21

- About Topology load balancing.....21
 - Understanding topology records.....21
 - Understanding how the BIG-IP system prioritizes topology records.....22
 - Creating a topology record.....24
 - Deleting a topology record.....25
- About Topology load balancing for a wide IP.....25
 - Example configuration: Topology load balancing for a wide IP.....25
 - Configuring a wide IP for Topology load balancing.....27
- About Topology load balancing for a pool.....28
 - Example configuration: Topology load balancing for a pool.....28
 - Configuring a pool for Topology load balancing.....30
- About Topology load balancing for both wide IPs and pools.....30
- About Topology load balancing for CNAME wide IPs and pools.....31
- About IP geolocation data.....31
 - About topology records and IP geolocation data.....31
 - Downloading and installing updates to the IP geolocation data.....31
 - Reloading default geolocation data using the Configuration utility.....32
 - Reloading default geolocation data using tmsh.....32

Preparing BIG-IP DNS for Static Persist Load Balancing for Transaction-Oriented

- Traffic.....35**
 - About Static Persist load balancing.....35
 - About configuring BIG-IP DNS to provide Static Persist load balancing35
 - Configuring CIDR options for Static Persist load balancing.....35

- Legal Notices.....37**
 - Legal notices.....37

About Global Server Load Balancing

Introducing BIG-IP DNS

BIG-IP[®] DNS (formerly GTM[™]) is a system that monitors the availability and performance of global resources and uses that information to manage network traffic patterns. BIG-IP DNS uses load balancing algorithms, topology-based routing, and iRules[®] to control and distribute traffic according to specific policies.

About global server load balancing

BIG-IP[®] DNS provides tiered global server load balancing (GSLB). BIG-IP DNS distributes DNS name resolution requests, first to the best available pool in a wide IP, and then to the best available virtual server within that pool. BIG-IP DNS selects the best available resource using either a static or a dynamic load balancing method. Using a static load balancing method, BIG-IP DNS selects a resource based on a pre-defined pattern. Using a dynamic load balancing method, BIG-IP DNS selects a resource based on current performance metrics collected by the `big3d` agents running in each data center.

Static load balancing methods

This table describes the static load balancing methods available in BIG-IP[®] DNS.

Name	Description	Recommended Use	Wide IP Load Balancing	Preferred Method	Alternate Method	Fallback Method
Drop Packet	BIG-IP DNS drops the DNS request.	Use Drop Packet for the Alternate load balancing method when you want to ensure that BIG-IP DNS does not offer in a response a virtual server that is potentially unavailable.	No	Yes	Yes	Yes
Fallback IP	BIG-IP DNS distributes DNS name resolution requests to a virtual server that you specify. This virtual server is not monitored for availability.	Use Fallback IP for the fallback load balancing method when you want BIG-IP DNS to return a disaster recovery site when the preferred and alternate load balancing methods do not return an available virtual server.	No	No	No	Yes

Name	Description	Recommended Use	Wide IP Load Balancing	Preferred Method	Alternate Method	Fallback Method
Global Availability	BIG-IP DNS distributes DNS name resolution requests to the first available virtual server in a pool. BIG-IP DNS starts at the top of a manually configured list of virtual servers and sends requests to the first available virtual server in the list. Only when the virtual server becomes unavailable does BIG-IP DNS send requests to the next virtual server in the list. Over time, the first virtual server in the list receives the most requests and the last virtual server in the list receives the least requests.	Use Global Availability when you have specific virtual servers that you want to handle most of the requests.	Yes	Yes	Yes	Yes
None	BIG-IP DNS distributes DNS name resolution requests skipping either the next available pool in a multiple pool configuration or the current load balancing method. If all pools are unavailable, BIG-IP DNS returns an aggregate of the IP addresses of all the virtual servers in the pool using BIND.	Use None for the alternate and fallback methods when you want to limit each pool to a single load balancing method. If the preferred load balancing method fails, BIG-IP DNS offers the next pool in a load balancing response.	No	No	Yes	Yes
Ratio	BIG-IP DNS distributes DNS name resolution requests among the virtual servers in a pool or among pools in a multiple pool configuration using <i>weighted round robin</i> , a load balancing pattern in which requests are distributed among several resources based on a priority level or weight assigned to each resource.	Use Ratio when you want to send twice as many connections to a fast server and half as many connections to a slow server.	Yes	Yes	Yes	Yes
Return to DNS	BIG-IP DNS immediately distributes DNS name resolution requests to an LDNS for resolution.	Use Return to DNS when you want to temporarily remove a pool from service. You can also use Return to DNS when you want to limit a pool in a single pool configuration to only one or two load balancing attempts.	No	Yes	Yes	Yes
Round Robin	BIG-IP DNS distributes DNS name resolution requests in a circular and sequential pattern among the virtual servers in a pool. Over time each virtual server receives an equal number of requests.	Use Round Robin when you want to distribute requests equally among all virtual servers in a pool.	Yes	Yes	Yes	Yes

Name	Description	Recommended Use	Wide IP Load Balancing	Preferred Method	Alternate Method	Fallback Method
Static Persist	BIG-IP DNS distributes DNS name resolution requests to the first available virtual server in a pool using the persist mask with the source IP address of the LDNS and a hash algorithm to determine the order of the virtual servers in the list. This hash algorithm orders the virtual servers in the list differently for each LDNS that is passing traffic to the system taking into account the specified CIDR of the LDNS. Each LDNS (and thus each client) generally resolves to the same virtual server; however, when the selected virtual server becomes unavailable, BIG-IP DNS sends requests to another virtual server until the original virtual server becomes available. Then BIG-IP DNS again resolves requests to that virtual server.	Use Static Persist when you want requests from a specific LDNS to resolve to a specific virtual server.	No	Yes	Yes	Yes
Topology	BIG-IP DNS distributes DNS name resolution requests using proximity-based load balancing. BIG-IP DNS determines the proximity of the resource by comparing location information derived from the DNS message to the topology records in a topology statement you have configured.	Use Topology when you want to send requests from a client in a particular geographic region to a data center or server located in that region.	Yes	Yes	Yes	Yes

Dynamic load balancing methods

This table describes the dynamic load balancing methods available in BIG-IP® DNS.

Name	Description	Wide IP load balancing	Preferred method	Alternate method	Fallback method
Completion Rate	BIG-IP® DNS distributes DNS name resolution requests to the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client's LDNS.	No	Yes	No	Yes
CPU	BIG-IP DNS distributes DNS name resolution requests to the virtual server that currently has the most CPU processing time available.	No	Yes	No	Yes
Hops	BIG-IP DNS distributes DNS name resolution requests to a virtual server in the data center that has the fewest router hops from the client's LDNS.	No	Yes	No	Yes

Name	Description	Wide IP load balancing	Preferred method	Alternate method	Fallback method
Kilobytes/Second	BIG-IP DNS uses the traceroute utility to track the number of router hops between a client's LDNS and each data center. BIG-IP DNS distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of kilobytes per second. Use Kilobytes/Second only with virtual servers for which BIG-IP DNS can collect the kilobytes per second metric.	No	Yes	No	Yes
Least Connections	BIG-IP DNS distributes DNS name resolution requests to virtual servers on BIG-IP® Local Traffic Manager™ (LTM®) that currently hosts the fewest connections. Use Least Connections only with LTM servers.	No	Yes	No	Yes
Packet Rate	BIG-IP DNS distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of packets per second.	No	Yes	Yes	Yes
Quality of Service	BIG-IP DNS distributes DNS name resolution requests to virtual servers based on a score assigned to each virtual server that is calculated from current performance metrics. Use Quality of Service only when you have configured BIG-IP DNS to calculate an overall score for each virtual server based on performance metrics.	No	Yes	No	Yes
Round Trip Time	BIG-IP DNS distributes DNS name resolution requests to the virtual server with the fastest measured round trip time between a data center and a client's LDNS.	No	Yes	No	Yes
Virtual Server Score	BIG-IP DNS distributes DNS name resolution requests to virtual servers on LTM based on a user-defined ranking. Use Virtual Server Score only with LTM systems on which you have assigned scores to each virtual server.	No	Yes	Yes	Yes
Virtual Server Capacity	BIG-IP DNS distributes DNS name resolution requests to virtual servers in a list that are weighted by the number of available virtual servers in the pool. Use Virtual Server Capacity for load balancing virtual servers managed by LTM Systems. BIG-IP DNS selects a virtual server that has the most available (UP) members. When selecting a virtual server from a wide IP pool and two or more virtual servers result in equal scores, BIG-IP DNS will return one of the equal scored virtual servers randomly.	No	Yes	Yes	Yes

About load balancing and resource availability

BIG-IP® DNS load balances DNS name resolution requests to resources based on availability. A resource is available when it meets one or more pre-defined requirements. BIG-IP DNS uses three methods to determine resource availability: a dependency on another resource, limit settings, or a set of values returned by a monitor. When BIG-IP DNS considers a resource unavailable, BIG-IP DNS attempts to select the next resource based on the current load balancing method.

About virtual server dependency

Within BIG-IP® DNS, you can configure a virtual server to be available based on the availability of other virtual servers.

Consider the fictional company SiteRequest. One of the servers, serverMain, at the Tokyo data center has two virtual servers: vsContact, which points to the contacts page of the web site, and vsMail, which points to the mail system. The vsMail virtual server is in the Dependency List of the vsContact virtual server. As a result, BIG-IP DNS considers the vsContact virtual server available only if the vsMail virtual server is also available.

Configuring virtual server availability to be dependent on the status of other virtual servers

Ensure that multiple virtual servers are configured on the server. Determine the virtual servers upon which you want the availability of a virtual server to be dependent.

Configure a virtual server to be available based on the availability of other virtual servers by configuring a **Dependency List** for the virtual server.

1. On the Main tab, click **DNS > GSLB > Servers**.
The Server List screen opens.
2. In the Server List, click a server name.
The server settings and values display.
3. On the menu bar, click **Virtual Servers**.
A list of the virtual servers configured on the server displays.
4. In the Virtual Servers list, click a virtual server name.
The virtual server settings and values display.
5. From the **Configuration** list, select **Advanced**.
Additional controls display on the screen.
6. In the Dependency List area, from the **Virtual Servers** list, select each virtual server on which you want the virtual server to be dependent, and then click **Add**.
The virtual servers display in the list as you add them.
7. Click **Finished**.

The virtual server is now available only when the virtual servers on the dependency list are also available.

Limit settings for resource availability

This table describes the limit settings BIG-IP® DNS uses to determine resource availability. A *limit setting* is a threshold for a statistic associated with a system.

Limit setting	Server-level	Pool-level	Virtual Server-level	BIG-IP Systems	Other Load Balancers	Hosts
Maximum allowable throughput in bits per second	Y	Y	Y	Y	Y	Y
Packets	Y	Y	Y	Y	Y	Y
Current connections	Y	Y	Y	Y	Y	Y
Connection	N	N	Y	Y	N	N
CPU	Y	N	N	N	Y	Y
Memory	Y	N	N	N	Y	Y

About wide IP-level load balancing

BIG-IP® DNS selects pools based on the order in which they are listed in a wide IP. When you organize pools in conjunction with the Global Availability, Ratio, Round Robin, and Topology load balancing methods, consider the order in which the pools are listed in the Pool List.

The Global Availability load balancing method instructs BIG-IP DNS to select the first pool in the wide IP pool list until it becomes unavailable, and then to select the next pool in the list until the first pool becomes available again. This ensures that the most robust pool receives DNS name resolution requests, while the other pools act as backups in case the primary pool becomes unavailable.

Note: A pool is available if it can provide an answer from any of the configured response methods: Preferred load balancing (LB), Alternate LB, Fallback LB, or Fallback IP.

About the Global Availability load balancing method

The *Global Availability* load balancing method distributes DNS name resolution requests based on the order of resources in a list. Using global availability, BIG-IP® DNS sends a request to the first available resource in a list. Only when a resource becomes unavailable does BIG-IP DNS send requests to the next resource in the list. You can only select the next resource in the list if the Fallback load balancing method is set to None. Over time, the first resource in the list receives the most requests and the last resource in the list receives the least requests.

Testing global server load balancing without verifying availability of virtual servers

You can configure BIG-IP DNS load balancing in a staging environment to load balance DNS name resolution requests to virtual servers without verifying the availability of the virtual servers.

1. On the Main tab, click **DNS > Settings > GSLB > Load Balancing**.
The Load Balancing configuration screen opens.
2. Deselect the **Verify Virtual Server Availability** check box.
3. Click **Update**.

About the Ratio load balancing method

The *Ratio* load balancing method distributes DNS name resolution requests among the virtual servers in a pool or among pools in a multiple pool configuration using *weighted round robin*, a load balancing pattern in which requests are distributed among several resources based on a priority level or weight assigned to each resource.

Using the Ratio method, you can configure BIG-IP® DNS to send twice as many connections to a fast, new server, and half as many connections to an older, slower server.

About wide IPs and weighting pools for the Ratio load balancing method

When you configure a wide IP to use the Ratio load balancing method, BIG-IP® DNS load balances DNS name resolution requests across the pools in the wide IP based on the weight assigned to each pool. BIG-IP DNS uses pool weight as a percentage of the total of the weights of all the pools in the wide IP to determine the frequency at which a pool receives connection requests.

Consider the fictional company SiteRequest, where the wide IP `www.siterequest.com` contains three pools, with the following weight assignments:

- Pool 1: weight 50
- Pool 2: weight 25
- Pool 3: weight 25

Each time BIG-IP DNS selects this wide IP, it load balances DNS name resolution requests across all three pools. Over time, the load balancing statistics for this wide IP appear as follows:

- Pool 1: selected 50 percent of the time
- Pool 2: selected 25 percent of the time
- Pool 3: selected 25 percent of the time

About pools and weighting pool members for the Ratio load balancing method

When you configure a pool to use the Ratio load balancing method, BIG-IP® DNS, formerly Global Traffic Manager™ (GTM™), load balances requests across the pool members based on the weight assigned to each pool member (virtual server). The system uses pool member weight as a percentage of the total of the weights of all the members assigned to the pool to determine the frequency at which a pool member receives connection requests.

Consider the fictional company SiteRequest, where the wide IP `www.siterequest.com` contains a pool named `poolMain`. This pool contains three members, with the following weight assignments:

- Virtual Server 1: weight 50
- Virtual Server 2: weight 25
- Virtual Server 3: weight 25

Each time BIG-IP DNS selects this pool, it load balances across all three members. Over time, the load balancing statistics for this pool appear as follows:

- Virtual Server 1: selected 50 percent of the time
- Virtual Server 2: selected 25 percent of the time
- Virtual Server 3: selected 25 percent of the time

About the Round Robin load balancing method

The *Round Robin* load balancing method distributes DNS name resolution requests in a circular and sequential pattern among the virtual servers in a pool. Over time, each virtual server receives an equal number of connections.

About pool-level load balancing

BIG-IP® DNS provides three tiers of pool-level load balancing to identify a virtual server to handle a DNS name resolution request.

Preferred Load Balancing Method

The first load balancing method BIG-IP DNS uses to return the IP address of a virtual server in response to a DNS name resolution request. The preferred method can be either static or dynamic.

Alternate Load Balancing Method

If the preferred load balancing method fails to return a valid resource in response to a DNS name resolution request, it is likely that BIG-IP DNS was unable to acquire the proper metrics to perform load balancing. The alternate method can be only static.

Fallback Load Balancing Method

If the alternate load balancing method fails to return a valid resource in response to a DNS name resolution request, BIG-IP DNS uses the fallback method. To ensure that BIG-IP DNS returns a response to a request, the fallback method ignores the availability status of a resource. The fallback method can be either static or dynamic.

When the fallback method is set to Return to DNS, the response from BIND is processed the same as any other DNS packet. If the response from BIND matches a wide IP, the load balancing algorithm may use a downed pool. If the downed pool is not configured to select another pool based on availability of the pool members, you can select a downed pool. For this scenario, BIND is not used a second time.

If all of the configured load balancing methods fail to provide a valid resource in response to a DNS name resolution request, either the request fails or BIG-IP DNS uses the local BIND to resolve the request.

About the Drop Packet load balancing method

The *Drop Packet* load balancing method indicates that BIG-IP® DNS drops a DNS name resolution request. This load balancing method is most often selected for the Alternate load balancing method to ensure that BIG-IP DNS does not return an IP address for an unavailable resource.

About the Virtual Server Score load balancing method

The *Virtual Server Score* load balancing method distributes DNS name resolution requests to pool members (virtual servers) based on a user-defined ranking system.

Note: This method can be used only for distributing requests to pool members controlled by BIG-IP® Local Traffic Manager™ (LTM™) systems.

About the Virtual Server Capacity load balancing method

The *Virtual Server Capacity* load balancing method distributes DNS name resolution requests to pool members (virtual servers) based on a system-generated list of pool members (virtual servers) weighted by capacity. BIG-IP DNS selects the pool member with the greatest capacity most often, but over time, all pool members are returned in responses. When pool members have the same capacity, BIG-IP DNS uses the Round Robin method to select a pool member.

About the Round Trip Times load balancing method

The *Round Trip Times* load balancing method distributes DNS name resolution requests to the pool member (virtual server) with the fastest measured round trip time between a data center and a client's LDNS.

About the Packet Rate load balancing method

The *Packet Rate* load balancing method distributes DNS name resolution requests to the pool member (virtual server) that is currently processing the fewest number of packets per second.

About the Least Connections load balancing method

The *Least Connections* load balancing method distributes DNS name resolution requests to pool members (virtual servers) that are managed by load balancing servers, such as BIG-IP® Local Traffic Manager™ (LTM™). BIG-IP DNS selects a pool member that currently hosts the fewest connections.

About the Kilobyte/Second load balancing method

The *Kilobyte/Second* load balancing method distributes DNS name resolution requests to the pool member (virtual server) that is currently processing the fewest number of kilobytes per second.

Note: This method can be used only with servers for which BIG-IP DNS can collect the kilobytes per second metric.

About the Hops load balancing method

The *Hops* load balancing method distributes DNS name resolution requests based on the `traceroute` utility and tracks the number of intermediate system transitions (router hops) between a client's LDNS and each data center. BIG-IP DNS distributes requests to a pool member in the data center that is the fewest router hops from the LDNS.

About the Completion Rate load balancing method

The *Completion Rate* load balancing method distributes DNS name resolution requests to the pool member (virtual server) that currently maintains the least number of dropped or timed-out packets during a transaction between a pool member in a data center and the client's LDNS.

About the CPU load balancing method

The *CPU* load balancing method distributes DNS name resolution requests to the pool member (virtual server) that currently has the most CPU processing time available.

About the Return to DNS load balancing method

The *Return to DNS* load balancing method immediately returns DNS name resolution requests to the LDNS for resolution. When you use this load balancing method, for client queries, the BIG-iP system increments the Return to DNS statistics; otherwise, the system increments the Return from DNS statistics.

Use this method when you want to temporarily remove a pool from service or when you want to limit a pool, in a single pool configuration, to only one or two request attempts.

About Static Persist load balancing

The Static Persist load balancing method uses the persist mask, with the source IP address of the LDNS, in a deterministic algorithm to send requests to a specific pool member (virtual server). Using this method, BIG-IP[®] DNS sends DNS name resolution requests to the first available pool member based on a hash algorithm that determines the order of the pool members. This algorithm orders the pool members differently for each LDNS that is sending requests to BIG-IP DNS, taking into account the Classless Inter-Domain Routing (CIDR) of the LDNS. As BIG-IP DNS distributes requests across all pool members, requests from each LDNS (and thus, each client) are generally sent to the same pool member. When the selected pool member becomes unavailable, BIG-IP DNS sends requests to another pool member. When the original pool member becomes available again, BIG-IP DNS sends requests to that pool member.

About the Fallback IP load balancing method

The *Fallback IP* load balancing method distributes DNS name resolution requests to a specific user-specified IP address. This IP address is not monitored for availability. Use this load balancing method only for the Fallback IP method and specifically to provide a disaster recovery site.

Verifying the availability of virtual servers when using the fallback load balancing method

You can configure BIG-IP DNS to verify that a virtual server is up before returning the IP address of the virtual server in a response to a DNS name resolution request. Do this when the preferred and alternate load balancing methods assigned to a pool do not return a valid response and BIG-IP DNS begins to use the configured fallback load balancing method.

1. On the Main tab, click **DNS > Settings > GSLB > Load Balancing**.
The Load Balancing configuration screen opens.
2. Select the **Respect Fallback Dependency** check box.
3. Click **Update**.

About the None load balancing method

The *None* load balancing method skips the current load balancing method, distributes DNS name resolution requests to the next available pool in a multi-pool configuration.

If the alternate load balancing method for a pool is None, BIG-IP DNS skips the alternate method and immediately tries the fallback method. If the fallback method is None, and there are multiple pools configured, BIG-IP DNS uses the next available pool. You can only select the next pool in the list if the fallback load balancing method is set to None. If all pools are unavailable, BIG-IP DNS returns an aggregate of the IP addresses of all pool members using BIND. Alternatively, when the preferred method for all pools is configured, but the alternate and fallback methods are set to None, if the preferred method fails, BIG-IP DNS uses the next available pool.

About the QoS load balancing method

The *Quality of Service (QoS)* dynamic load balancing method uses current performance metrics to calculate an overall QoS score for each pool member (virtual server). When load balancing DNS name resolution requests, BIG-IP DNS selects a virtual server with the best overall QoS score. If virtual servers have identical scores, BIG-IP® DNS load balances connections to those virtual servers using the round robin method. If QoS scores cannot be determined, BIG-IP DNS load balances connections across all pool members using the round robin method.

Understanding the QoS equation

The equation for calculating the overall Quality of Service (QoS) score is:

```
POOL_CONFIG->rtt * (GLOBALS->rtt / path->rtt) * 10 +
POOL_CONFIG->hops * (GLOBALS->hops / path->hops) +
POOL_CONFIG->hit_ratio * (path->hit_ratio / GLOBALS->hit_ratio)
```

```

POOL_CONFIG->packet_rate * GLOBALS->packet_rate / (vs->packet_rate * 100) +
POOL_CONFIG->bps * (GLOBALS->bps / vs->bps) +
POOL_CONFIG->topology * (topology_match->score / GLOBALS->topology) +
POOL_CONFIG->vs_capacity * vs->cur_serv_cnt +
POOL_CONFIG->vs_score * vs->cur_vs_score +
POOL_CONFIG->lcs * vs->link->lcs * 10
    
```

Pool members (virtual servers) inherit the QoS settings from the pool. In the equation, the value of POOL_CONFIG->"setting name" can be found in the properties of a pool, the value of GLOBALS->"setting name" in the global BIG-IP DNS setting, and the value of path->"setting name" These are measured values that come from path metrics. If there are no path metrics, the system does not perform path metric calculations and computes the QoS score using the other calculations. vs->"field" These are measured values that come from measurements the system makes on virtual servers. If there are no measurements, the system does not perform these calculations and computes the QoS score using the other calculations. Each QoS coefficient, its scale, default value, upper limit, and whether a higher or lower value is more efficient are defined in the table.

Table 1: QoS coefficients defined

Coefficient	Scale	Default value	Upper limit	Is higher or lower value more efficient?
Round trip time (rtt)	Microseconds	50	2,000,000	L
Completion rate (hit ratio)	Percentage of successfully transferred packets (0-100%)	5	100%	H
Hops	Number of intermediate systems transitions	0	64	L
Packet rate	Packets per second	1	700	L
bits/second	Bits per second throughput	3	15000	L
Topology	Score that defines network proximity by comparing server and LDNS IP addresses (0-2 ³²)	0	100	H
Virtual server capacity (vs capacity)	Number of nodes up	0	20	H
Virtual server score (vs score)	User-defined ranking of virtual servers	0	100	H
Link capacity (lcs)	Based on the target dynamic ratio	30	2,000,000	H

About customizing the QoS equation

When you customize the QoS equation, consider these three concepts:

Scale

The raw metrics for the coefficients in the QoS equation are on different scales. For example, completion rate is measured in percentages, while packet rate is measured in packets per second.

Normalization

BIG-IP DNS normalizes the raw metrics to values in the range of 0 - 10.

Emphasis

You can adjust coefficients to emphasize one normalized metric over another.

When you customize the QoS equation configuration using the values in the table, if the completion rates for two virtual servers are close, the system chooses the virtual server with the best packet rate. If both the completion rates and the packet rates are close, the round trip time (RTT) breaks the tie. In this example, BIG-IP DNS does not use the metrics for topology, hops, link capacity, vs capacity, and kilobytes/second to determine how to distribute connections.

Note: You can set a value for either RTT or hops. If you set both, BIG-IP DNS incorporates the RTT and resets the hops to 0 (zero).

Coefficient	Value
Round Trip Time	50
Hops	0
Topology	0
Completion Rate	5
Packet Rate	10
VS Capacity	0
Bits/second	35
Link Capacity	30
Virtual Server Score	10
Kilobytes/Second (KBPS)	3

Customizing the QoS equation for load balancing global traffic

Determine the pool to which you want to apply a customized QoS equation.

Customize the QoS equation to load balance the DNS name resolution requests the members of this pool handle.

1. On the Main tab, click **DNS > GSLB > Pools**.
2. Click the name of the pool for which you want to modify the QoS equation.
The Pool Properties screen displays.
3. On the menu bar, click **Members**.
The Members Properties screen displays.
4. Select **Quality of Service** from either the **Preferred** or **Fallback** list.
The Quality of Service Weights area displays.
5. Define the QoS coefficients for this pool.
6. Click **Update**.

About dynamic ratio load balancing

When you use dynamic ratio load balancing, BIG-IP DNS treats dynamic load balancing values as ratios, and distributes DNS name resolution requests to the virtual servers in the pool in proportion to these ratios.

Consider a pool named `primaryOne`, which contains two virtual servers: `memberOne` and `memberTwo`. `primaryOne` is configured with the Preferred load balancing method set to **Round Trip Time**. BIG-IP DNS determines that the round trip time for `memberOne` is 50 microseconds and the round trip time for `memberTwo` is 100 microseconds. When the **Dynamic Ratio** setting on the `primaryOne` pool is disabled, BIG-IP DNS always sends DNS name resolution requests to `memberOne`, because that virtual server has the lowest round trip time value. When the **Dynamic Ratio** setting on the `primaryOne` pool is enabled, BIG-IP DNS treats the round trip time values as ratios and sends twice as many DNS name resolution requests to `memberOne` as it sends to `memberTwo`, because the round trip time for `memberOne` is twice as fast as the round trip time for `memberTwo`.

Distributing DNS requests based on weighted virtual servers

Determine the pool to which you want to apply the dynamic ratio feature.

Configure BIG-IP DNS to use dynamic load balancing values as ratios, and distribute DNS name resolution requests to virtual servers in a pool in proportion to these ratios.

1. On the Main tab, click **DNS > GSLB > Pools**.
The Pools list screen opens.
2. Click the name of the pool that you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. Select the **Dynamic Ratio** check box.
5. Click **Update**.

Using the preferred load balancing method when metrics are unavailable

Configure BIG-IP DNS to use the preferred load balancing method assigned to a pool even when metrics for the pool are unavailable. BIG-IP DNS uses old metrics, rather than the alternate load balancing method assigned to the pool.

1. On the Main tab, click **DNS > Settings > GSLB > Load Balancing**.
The Load Balancing configuration screen opens.
2. Select the **Ignore Path TTL** check box.
3. Click **Update**.

BIG-IP DNS uses path information gathered during metrics collection even if the time-to-live (TTL) value of that information has expired.

Configuring the resources in a pool for manual resume

Determine the pool to which you want to apply the manual resume feature.

When a virtual server goes offline, BIG-IP DNS proceeds to send DNS name resolution requests to other virtual servers, based on the current load balancing method. By default, when the virtual server becomes

available again, BIG-IP DNS resumes sending requests to that resource. When you do not want BIG-IP DNS to resume to send requests to the virtual servers in a pool immediately after the resources become available, enable the manual resume feature on the pool.

1. On the Main tab, click **DNS > GSLB > Pools**.
The Pools list screen opens.
2. Click the name of the pool that you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. Select the **Manual Resume** check box.
5. Click **Update**.

After a virtual server in this pool goes offline, you must manually enable the virtual server before BIG-IP DNS can resume sending requests to the virtual server.

Restoring availability of a pool member manually

Determine the virtual server that you want to manually enable.

When a virtual server in a pool that is configured for manual resume becomes available, you must manually enable the virtual server before BIG-IP DNS can begin sending DNS name resolution requests to the virtual server.

1. On the Main tab, click **DNS > GSLB > Pools**.
The Pools list screen opens.
2. Click the name of the pool to which the virtual server you want to enable belongs.
3. On the menu bar, click **Members**.
4. Select the check box next to the virtual server that you want to enable, and then click **Enable**.

The virtual server is now available to receive DNS name resolution requests.

Best Practices: Dynamic load balancing and performance configuration

How do I ensure that the BIG-IP® DNS (formerly GTM) has timely access to path and network traffic metrics?

Run a `big3d` agent on at least one BIG-IP® system in each data center to ensure that the BIG-IP DNS has timely access to path and network traffic metrics.

How do I ensure that BIG-IP DNS can gather information for dynamic load balancing?

If you use dynamic load balancing modes, you must run a `big3d` agent on every BIG-IP system in your network.

How do I configure BIG-IP DNS for the best performance?

The load on the `big3d` agents depends on the timer settings that you assign to the different types of data the agents collect. The shorter the timers, the more frequently the agent needs to refresh the data. While short timers guarantee that you always have valid data readily available for load balancing, they also increase the frequency of data collection.

About Global Server Load Balancing

The more local DNS servers that make resolution requests, the more path data that the `big3d` agents have to collect. While round trip time for a given path may vary constantly due to current network load, the number of hops along a network path between a data center and a specific LDNS does not often change. Consequently, you may want to set short timer settings for round trip time data so that it refreshes more often, but set high timer settings for hops data because it does not need to be refreshed often.

Troubleshooting

If you turn off the `big3d` agent on a BIG-IP system, the BIG-IP DNS can no longer check the availability of the server or its virtual servers. Therefore, the statistics screens display the status of these servers as unknown.

Using Topology Load Balancing to Distribute DNS Requests to Specific Resources

How do I configure BIG-IP DNS to load balance DNS requests to specific resources?

You can configure BIG-IP® DNS to load balance DNS requests to a resource based on the physical proximity of the resource to the client making the request. You can also configure BIG-IP DNS to deliver region-specific content, such as news and weather, to a client making a request from a specific location.

You can accomplish this by configuring BIG-IP DNS to perform Topology load balancing.

About Topology load balancing

Topology load balancing distributes DNS name resolution requests based on the proximity of the client to the data center housing the resource that responds to the request. When Topology load balancing is enabled, the BIG-IP® system uses topology records to make load balancing decisions.

Understanding topology records

A *topology record* is a set of characteristics that maps the origin of a DNS name resolution request to a destination. Each topology record contains the following elements:

- A *request source statement* that specifies the origin LDNS of a DNS request.
- A *destination statement* that specifies the pool or pool member to which the weight of the topology record will be assigned.
- A weight that the BIG-IP® system assigns to a pool or a pool member during the load balancing process.

Note: In *tmsih*, the weight parameter is called *score*.

Understanding user-defined regions

A *region* is a customized collection of topologies that defines a specific geographical location that has meaning for your network. For example, you can create two custom regions named `Region_east` and `Region_west`. `Region_east` includes the states on the east coast of the United States. `Region_west` includes the states on the west coast of the United States. Then, you can use those custom regions as the Request Source or Destination of a topology record you create.

This table describes how the use of topology regions improves the load-balancing performance of the BIG-IP® system.

Faster load balancing configuration	Slower load balancing configuration
2 data centers	2 data centers
1000 pool members in each data center	1000 pool members in each data center
2 regions with 5000 CIDR entries each	
2 topology records:	10,000 topology records:
1 entry routes all requests from Region_east to data center1	5000 CIDR topology records route requests to data center1
1 entry routes all requests from Region_west to data center2	5000 CIDR topology records route requests to data center2

Creating a region for Topology load balancing

Create regions to customize the Topology load balancing capabilities of the BIG-IP® system. For example, you can create two regions to represent the data centers in your network: dc1_pools and dc2_pools. Alternatively, you can create a region to which you can add IP subnets as you expand your network. Then, when you create a topology record, you can use the custom regions as the Request Source or Destination of the record.

1. On the Main tab, click **DNS > GSLB > Topology > Regions**.
2. Click **Create**.
The new record screen opens.
3. In the **Name** field, type a unique identifier for the region.
4. To add members to the region, do the following for each member you want to add to the region:
 - a) From the **Member Type** list, select a type of identifier.
 - b) Select an operator, either **is** or **is not**.
 - c) From the **Continent** list, select the continent that contains the locations in the region you are creating.
 - d) Click **Add**.
5. Click **Create**.

You can now create a topology record using the custom region you created.

Understanding how the BIG-IP system prioritizes topology records

When Topology load balancing is configured, the order of the topology records is vital and affects how the BIG-IP® system scores the pools or pool members to which it load balances DNS name resolution requests. By default, the BIG-IP system prioritizes topology records using Longest Match sorting. As a result, topology records are automatically sorted based on a specific criteria each time the BIG-IP system configuration loads. Alternatively, you can disable Longest Match sorting and customize the order of the topology records in the list.

Understanding Longest Match topology record sorting

When Longest Match is enabled, the BIG-IP® system sorts the topology records by the LDNS request source statement, the destination statement, and the weight of the record.

The system first sorts the topology records by the type of LDNS request source statement using this order from highest to lowest:

1. IP subnet in CIDR format (the system places the most specific IP subnet at the top of the list; for example, 10.15.1.1/32, 10.15.1.0/24, 10.15.0.0/16, 10.0.0.0/8)
2. Region
3. ISP
4. State
5. Country
6. Continent
7. LDNS Request Source negation (record that excludes an LDNS)
8. Wildcard record (the system sorts the wildcard record to the bottom of the list, because this record is the least specific)

If the type of LDNS request source statement is the same in multiple topology records, the BIG-IP system then sorts these records by the type of destination statement using this order from highest to lowest:

1. IP subnet in CIDR format (the system places the most specific IP subnet at the top of the list; for example, 10.15.1.1/32, 10.15.1.0/24, 10.15.0.0/16, 10.0.0.0/8)
2. Data center
3. Pool
4. Region (customized collection of criteria)
5. ISP
6. State
7. Country
8. Continent
9. Destination negation (record that excludes a destination)
10. Wildcard record (the system sorts the wildcard record to the bottom of the list, because this record is the least specific)

If the type of LDNS request source statement is the same in multiple topology records and the type of destination statement is the same in those records, the system then uses the value of the weight from highest to lowest to sort the records.

The example shows a list of topology records sorted automatically using Longest Match. Note that the fourth and fifth records have the same LDNS subnet and the destinations are both of type State. Therefore, the weight determines the position in the list; thus, the record with the highest weight is first.

1. Idns: subnet 192.168.69.133/32 destination: subnet 10.15.1.1/32 weight: 500
2. Idns: subnet 192.168.69.133/32 destination: datacenter /Common/NorthAmerica weight: 400
3. Idns: subnet 192.168.69.0/24 destination: pool /Common/NorthAmerica weight 300
4. Idns: subnet 192.168.0.0/16 destination: state NY weight 200
5. Idns: subnet 192.168.0.0/16 destination: state WA weight 100

Customizing the sort order of topology records

Determine the order in which you want the topology records you create to be sorted.

Change the sort order of the topology records when you do not want the system to use the Longest Match sort order.

1. On the Main tab, click **DNS > GSLB > Topology > Records**.
2. Click the **Change Order** button.
3. Clear the **Longest Match** check box.

4. To change the order of the records in the Topology Record List, do the following:
 - a) From the list, select a topology record.
 - b) Click the **Up** or **Down** button to move the record to the preferred position in the list.
5. Click **Update**.

The BIG-IP system uses the customized Topology Record List for topology load balancing.

Important: *The BIG-IP system saves only one set of ordered topology records; if you re-enable Longest Match, your custom ordering will no longer be available.*

Configuring Longest Match

Ensure that topology records exist in the configuration.

Configure the BIG-IP[®] system to order the topology records using Longest Match.

1. On the Main tab, click **DNS > GSLB > Topology > Records**.
2. Click the **Change Order** button.
3. Select the **Longest Match** check box.
4. Click **Update**.

The BIG-IP system uses Longest Match sorting to order the topology records in a list.

Creating a topology record

Before you create topology records, it is essential that you understand how the system sorts the topology record list. Additionally, you must understand how the system uses the ordered list of records to assign scores to the pools or pool members, to which the BIG-IP[®] system load balances DNS requests.

Create topology records that instruct the BIG-IP system where to route DNS name resolution requests when Topology load balancing is enabled.

Tip: *The BIG-IP system is more efficient when using regions for Topology load balancing.*

1. On the Main tab, click **DNS > GSLB > Topology**.
2. Click **Create**.

The new record screen opens.
3. To create an LDNS request source statement, use the **Request Source** settings:
 - a) Select an origin type from the first list.
 - b) Select an operator, either **is** or **is not**.
 - c) Define the criteria for the request source statement based on the request source type you selected.
4. To create a destination (server object) statement, use the **Destination** settings:
 - a) Select a destination type from the first list.
 - b) Select an operator, either **is** or **is not**.
 - c) Define the criteria for the destination statement based on the destination type you selected.
5. In the **Weight** field, specify the priority of this record.
6. Click **Create**.

Deleting a topology record

Delete existing topology records as your network changes. For example, when you add a new data center to your network, the topology records that the BIG-IP® system uses to distribute DNS name resolution requests can become obsolete, requiring deletion.

Note: You cannot modify topology records; you can delete records and create new ones that meet your needs.

1. On the Main tab, click **DNS > GSLB > Topology**.
2. Select the topology record that you want to remove from the topology records list by selecting the corresponding **Select** check box.
3. Click **Delete**.
A confirmation screen appears.
4. Click **Delete**.

About Topology load balancing for a wide IP

When you use the topology load balancing method at a wide IP level with topology records that have a Data Center destination, the topology records have no effect. (This is because load balancing at a wide IP level selects between GTM pools, and GTM pools do not have a data center associated with them.) Topology records that have a Data Center destination have an effect only when using the topology load balancing method at the pool level.

Example configuration: Topology load balancing for a wide IP

This example illustrates how DNS name resolution requests are load balanced when a wide IP is configured for Topology load balancing. An administrator configures the wide IP `www.siterequest.net` for Topology load balancing. The wide IP contains three pools: Pool1 and Pool3 are located in the North America data center; Pool2 is located in the South America data center. Next, the administrator creates topology records, as shown in this figure, and ensures that Longest Match is enabled on the BIG-IP® system.

<input checked="" type="checkbox"/>	Order	LDNS Request Source	Destination	Weight
<input type="checkbox"/>	1	IP Subnet is 11.1.0.0/16	Pool is Pool1	100
<input type="checkbox"/>	2	IP Subnet is 10.1.0.0/16	Pool is Pool2	100
<input type="checkbox"/>	3	IP Subnet is 0.0.0.0/0	Pool is Pool3	10

Figure 1: Topology records for a wide IP configured for Topology load balancing

The first topology record directs all DNS name resolution requests from an LDNS in the IP subnet `11.1.0.0/16` to Pool1. The second topology record directs all DNS name resolution requests from an LDNS in the IP subnet `10.1.0.0/16` to Pool2. The third topology record is least specific. It directs DNS

name resolution requests from an LDNS in any IP subnet to Pool3. However, it is important to note that the weight of the third topology record is lower than the weights of the other topology records.

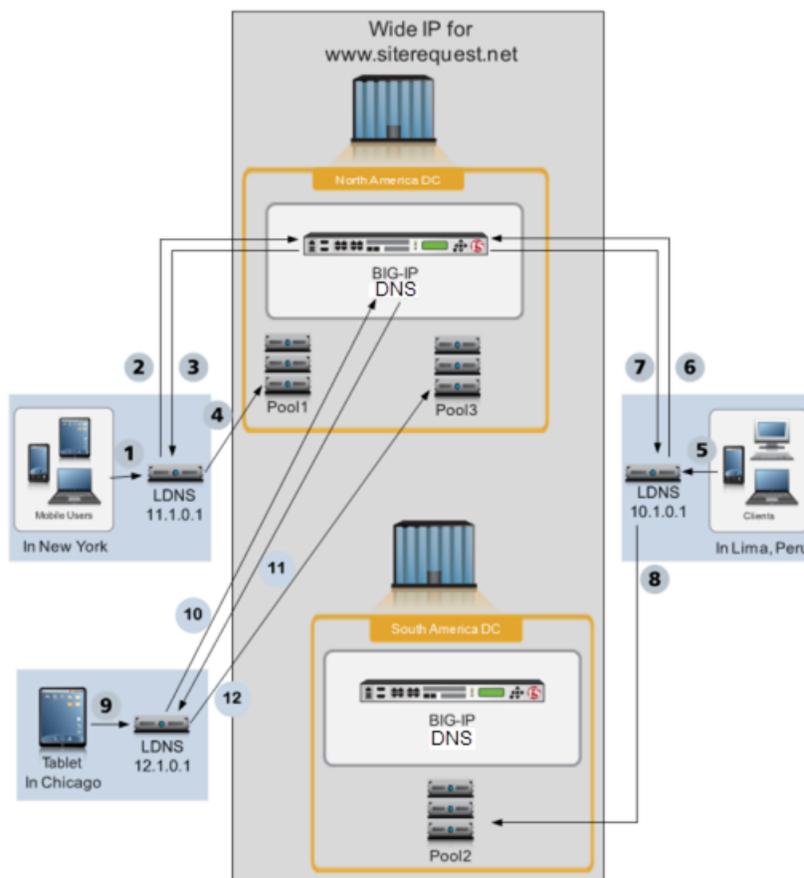


Figure 2: BIG-IP system load balancing DNS requests using a wide IP configured for Topology load balancing

1. A client in New York makes a DNS request.
2. LDNS 11.1.0.1 queries the BIG-IP system in the North America data center.
3. The BIG-IP system directs the LDNS to Pool1. To determine this answer, for each pool, one at a time, the BIG-IP system iterates through the list of two topology records to find a match. Pool1 matches the first topology record in the list, because both the LDNS request source (11.1.0.1) and the Destination (Pool1) of the DNS request match the first topology record; therefore, the BIG-IP system assigns a score of 100 to Pool1. For Pool2, there is no matching topology record that contains both the LDNS request source (11.1.0.1) and the Destination (Pool2); therefore, the BIG-IP system assigns a score of zero to Pool2. Pool3, matches the third topology record in the list, because both the LDNS request source (11.1.0.1) and the Destination (Pool3) of the DNS request match the third topology record; therefore, the BIG-IP system assigns a score of 10 to Pool3. The BIG-IP system directs the LDNS to send the request to the pool with the highest score.
4. The LDNS sends the DNS request to Pool1 in the North America data center. How the system distributes the DNS requests to the members of Pool1 is not depicted in this illustration, but is based on the load balancing method configured for Pool1.
5. A client in Lima makes a DNS request.
6. LDNS 10.1.0.1 queries the BIG-IP system in the North America data center.
7. The BIG-IP system directs the LDNS to Pool2. To determine this answer, for each pool, one at a time, the BIG-IP system iterates through the list of two topology records to find a match. For Pool1, there is not a matching topology record that contains both the LDNS request source (10.1.0.1) and the Destination

(Pool1); therefore, the BIG-IP system assigns a score of zero to Pool1. Pool2 matches the second topology record in the list, because both the LDNS request source (10.1.0.1) and the Destination (Pool2) of the DNS request match the second topology record; therefore, the BIG-IP system assigns a score of 100 to Pool2. Pool3, matches the third topology record in the list, because both the LDNS request source (10.1.0.1) and the Destination (Pool3) of the DNS request match the third topology record; therefore, the BIG-IP system assigns a score of 10 to Pool3. The BIG-IP system directs the LDNS to send the request to the pool with the highest score.

8. The LDNS sends the DNS request to Pool2 in the South America data center. How the system distributes the DNS requests to the members of Pool2 is not shown in this illustration, but is based on the load balancing method configured for Pool2.
9. A client in Chicago makes a DNS request.
10. LDNS 12.1.0.1 queries the BIG-IP system in the North America data center.
11. The BIG-IP system directs the LDNS to Pool3. To determine this answer, for each pool, one at a time, the BIG-IP system iterates through the list of two topology records to find a match. For Pool1, there is not a matching topology record that contains both the LDNS request source (12.1.0.1) and the Destination (Pool1); therefore, the BIG-IP system assigns a score of zero to Pool1. For Pool2, there is not a matching topology record that contains both the LDNS request source (12.1.0.1) and the Destination (Pool1); therefore, the BIG-IP system assigns a score of zero to Pool2. Pool3, matches the third topology record in the list, because both the LDNS request source (12.1.0.1) and the Destination (Pool3) of the DNS request match the third topology record; therefore, the BIG-IP system assigns a score of 10 to Pool3. The BIG-IP system directs the LDNS to send the request to the pool with the highest score.
12. The LDNS sends the DNS request to Pool3 in the North America data center. How the system distributes the DNS requests to the members of Pool3 is not depicted in this illustration, but is based on the load balancing method configured for Pool3.

Configuring a wide IP for Topology load balancing

Before you configure a wide IP for Topology load balancing, ensure the following:

- At least two pools are associated with the wide IP that you are configuring for Topology load balancing.
- Topology records that define how you want the BIG-IP® system to load balance DNS name resolution requests are configured.

You can use Topology load balancing to distribute DNS name resolution requests among the pools in a wide IP based on the geographic location of both the client making the request and the pool that handles the response.

1. On the Main tab, click **DNS > GSLB > Wide IPs**.
The Wide IP List screen opens.
2. Click the name of the wide IP you want to modify.
3. On the menu bar, click **Pools**.
4. From the **Load Balancing Method** list, select **Topology**.
5. Click **Update**.

Repeat this process for each wide IP that you want to configure for Topology load balancing.

About Topology load balancing for a pool

When you configure a pool for Topology load balancing, you can route DNS requests to the data center that is closest to the client making the request. With this configuration, the BIG-IP® system load balances DNS name resolution requests to the members of the pool.

Example configuration: Topology load balancing for a pool

This example illustrates how DNS name resolution requests are load balanced when a pool is configured for Topology load balancing. An administrator configures pools in two different data centers: the North America data center (North America DC) and the South America data center (South America DC) for Topology load balancing. A server that contains the pool members 10.10.10.1 - 10.10.10.3 resides in the North America DC. The server that contains the pool members 11.10.10.1 - 11.10.10.3 resides in the South America DC. Next, the administrator creates topology records, as shown in the following figure, to load balance DNS requests to members of the pools, and ensures that Longest Match is enabled on the BIG-IP® system.

Order	LDNS Request Source	Destination	Weight
1	Country is Bolivia	Data Center is SouthAmericaDC	10
2	Country is Peru	Data Center is SouthAmericaDC	10
3	Country is United States	Data Center is NorthAmericaDC	20
4	Country is Canada	Data Center is NorthAmericaDC	20

Figure 3: Topology records for a pool configured for Topology load balancing

The first topology record directs all DNS name resolution requests from an LDNS in Bolivia to the South America DC. The second topology record directs all DNS name resolution requests from an LDNS in Peru to the South America DC. The third topology record directs all DNS name resolution requests from an LDNS in the United States to the North America DC. The fourth topology record directs all DNS name resolution requests from an LDNS in Canada to the North America DC.

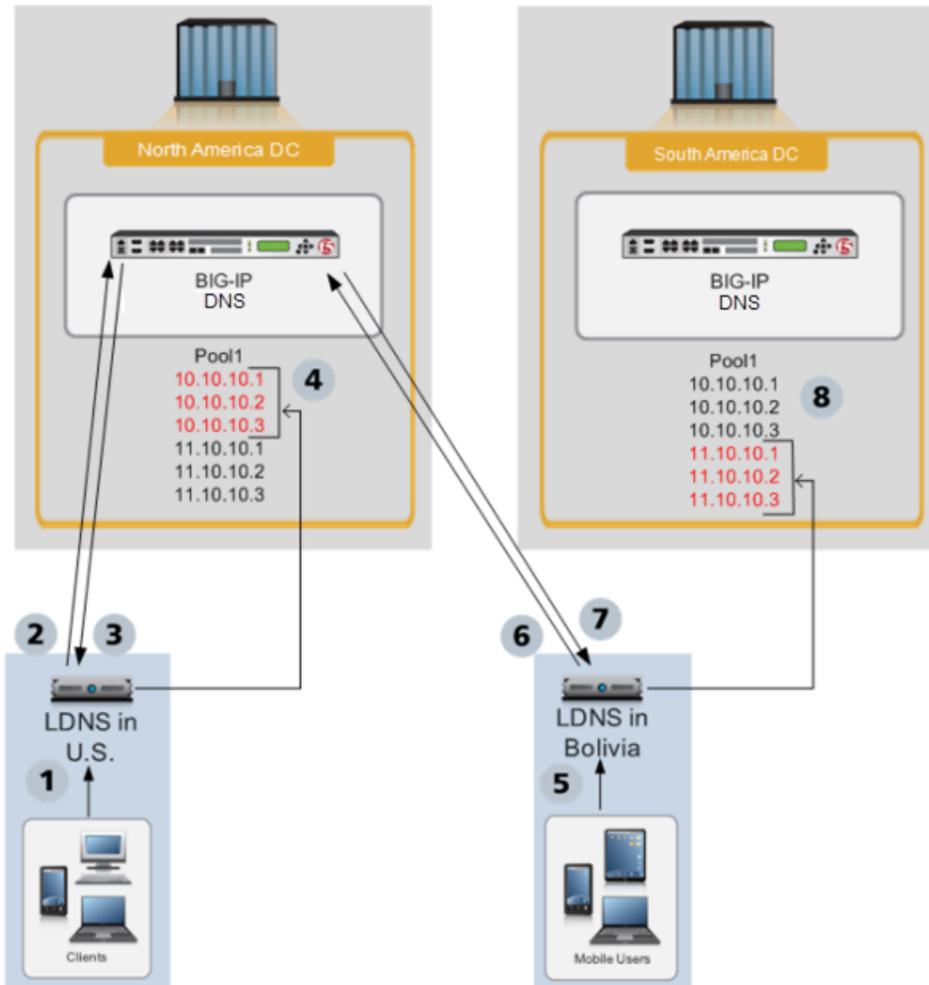


Figure 4: Pool configured for Topology load balancing

1. A client in the U.S. makes a DNS request.
2. An LDNS in the U.S. queries the BIG-IP system in the North America DC.
3. The BIG-IP system directs the LDNS to a member of Pool1 in the North America DC. To determine this answer, for each pool member, one at a time, the BIG-IP system iterates through the list of topology records to find a match. Pool members 10.10.10.1 - 10.10.10.3 each match the third topology record in the list, because both the LDNS request source (U.S.) and the Destination (North America DC) of the DNS request match the third topology record; therefore, the BIG-IP system assigns a score of 20 to each of those pool members. For each of the pool members 11.10.10.1 - 11.10.10.3, there is no matching topology record that contains both the LDNS request source (U.S.) and the Destination (South America DC); therefore, the BIG-IP system assigns a score of zero to each of those pool members. The BIG-IP system directs the LDNS to send the request to the pool member with the highest score.
4. The LDNS sends the DNS request to a pool member in the North America DC. Because all of the pool members in the North America DC have the same score, the system distributes the DNS requests to the pool members in a round robin fashion.
5. A client in Bolivia makes a DNS request.
6. An LDNS in Bolivia queries the BIG-IP system in the North America DC.
7. The BIG-IP system directs the LDNS to a pool member in the South America DC. To determine this answer, for each pool member, one at a time, the BIG-IP system iterates through the list of topology records to find a match. For each of the pool members 10.10.10.1 - 10.10.10.3 there is no matching

topology record that contains both the LDNS request source (Bolivia) and the Destination (North America DC); therefore, the BIG-IP system assigns a score of zero to each of those pool members. Pool members 11.10.10.1 - 11.10.10.3 each match the first topology record in the list, because both the LDNS request source (Bolivia) and the Destination (South America DC) of the DNS request match the first topology record; therefore, the BIG-IP system assigns a score of 10 to each of those pool members. The BIG-IP system directs the LDNS to send the request to the pool member with the highest score.

8. The LDNS sends the DNS request to a pool member in the South America DC. Because all of the pool members in the South America DC have the same score, the system distributes the DNS requests to the pool members in a round robin fashion.

Configuring a pool for Topology load balancing

Before you configure a pool for Topology load balancing, ensure the following:

- The pool you are configuring for Topology load balancing contains at least two pool members.
- Topology records that define how you want the BIG-IP® system to load balance DNS name resolution requests are configured.

You can use Topology load balancing to distribute DNS name resolution requests among the members of a pool based on the geographic location of both the client making the request and the member of the pool that handles the response.

1. On the Main tab, click **DNS > GSLB > Pools**.
The Pools list screen opens.
2. Click the name of the pool you want to modify.
3. On the menu bar, click **Members**.
4. In the Load Balancing Method area, from the **Preferred** list, select **Topology**.
5. In the Load Balancing Method area, from the **Alternate** list, select **Round Robin**.
6. In the Load Balancing Method area, from the **Fallback** list, select **None**.
7. Click **Update**.

Repeat this process for each pool that you want to configure for Topology load balancing.

About Topology load balancing for both wide IPs and pools

You can configure a wide IP for Topology load balancing. You can also configure each pool in the wide IP for Topology load balancing. When you configure both a wide IP and the pools in the wide IP for Topology load balancing, the BIG-IP® system uses topology records to load balance DNS name resolution requests first to a pool in the wide IP, and then, to a member of the pool.

Note:

When configuring both the wide IP and the pools in the wide IP for Topology load balancing, it is important to set the **Fallback** load balancing method for each pool to **None**. If you do not, DNS can send a DNS request to a pool in the wide IP even when no pool members are available. In this case, the load balancing algorithm for the pool would then fall back to BIND (static DNS). When you set the **Fallback** load balancing method for each pool to **None**, if no members of a pool are available, BIG-IP DNS sends the DNS request to another pool in the wide IP.

About Topology load balancing for CNAME wide IPs and pools

For a CNAME query against a CNAME wide IP with a CNAME pool, you can use Topology load balancing at the wide IP level to make a pool selection, given that matching topology entries are configured to specify a CNAME pool on the wide IP. However, when using Topology load balancing at the pool level to make pool member selections, the BIG-IP® system relies on a metrics pull-up to get the topology scores needed to pick a pool member because the pool members are non-terminal.

If the pool members are terminal members (for example, on an A or AAAA type pool), then you can use them to match against the topology entries and get scores. But Topology load balancing cannot be used to get scores for non-terminal members without doing a metrics pull-up because there is no way to specify a wide IP or DNS name in a given topology entry. Therefore, for a CNAME query against a CNAME wide IP with a CNAME pool, you can use Topology load balancing to pick a pool at the wide IP level, but it will not be used (even if configured) at the pool level to pick a pool member because the BIG-IP system will not perform a metrics pull-up.

About IP geolocation data

The BIG-IP® system uses an IP geolocation database to determine the origin of DNS requests. The database included with the BIG-IP system provides geolocation data for IPv6 addresses at the continent and country levels. It also provides geolocation data for IPv4 addresses at the continent, country, state, ISP, and organization levels. The state-level data is worldwide, and thus includes designations in other countries that correspond to the U.S. state-level in the geolocation hierarchy, such as, provinces in Canada.

***Note:** If you require geolocation data at the city-level, contact your F5® Networks sales representative to purchase additional database files.*

About topology records and IP geolocation data

The BIG-IP® system uses an IP geolocation database to determine the IP addresses that match the geographic names that you define in a topology record, such as continent and country.

Downloading and installing updates to the IP geolocation data

You can download a monthly update to the IP geolocation database from F5® Networks. The BIG-IP® system uses the IP geolocation database to determine the origin of DNS name resolution requests.

1. Log in to the F5 Networks customer web site at <http://downloads.f5.com>, and click **Find a Download**.
2. In the F5 Product Family column, find **BIG-IP**, and then in the Product Line column, click either **BIG-IP v11.x/Virtual Edition**.
3. Select a version from the list preceding the table.
4. In the Name column, click **GeolocationUpdates**.
5. Click **I Accept** to accept the license.
6. In the Filename column, click the name of the most recent compressed file that you want to download.

7. In the Ready to Download table, click the download method that you want to use.
8. In the dialog box, click **OK**.
9. Select the directory in which you want to save the compressed file, and then decompress the file to save the RPM files on the system.
10. To install and load one of the RPM files, run this command (the path and file name are case-sensitive):

```
geoip_update_data -f </path to RPM file and file name >.
```

The system installs and loads the specified database file.
11. Repeat step 10 for each of the RPM files that you saved to the system in step 9.

You can access the ISP and organization-level geolocation data for IPv4 addresses only using the iRules[®] `whereis` command.

Reloading default geolocation data using the Configuration utility

Before you reload the default geolocation data, delete the RPM files that are in the `/shared/GeoIP` directory.

To uninstall an update to the IP geolocation database, reload the default geolocation database files using the Configuration utility.

1. At the BASH prompt, run this command to query the RPM database and determine what geolocation data is installed:

```
rpm -qa --dbpath /shared/lib/rpm/
```

The system returns a list of RPMs, for example:

```
geoip-data-ISP-1.0.0-20110203.61.0
geoip-data-Region2-1.0.0-20110203.61.0
geoip-data-Org-1.0.0-20110203.61.0
```

2. To uninstall the RPMs, run this command for each RPM in the list:

```
rpm -e --dbpath /shared/lib/rpm/ <name of file>
```

For example, to uninstall `geoip-data-ISP-1.0.0-20110203.61.0`, run this command: `rpm -e --dbpath /shared/lib/rpm/ geoip-data-ISP-1.0.0-20110203.61.0`

3. To remove the symlink in the `/shared/GeoIP` directory, run this command:

```
rm -f /shared/GeoIP/*
```

4. Log on to the Configuration utility.
5. On the Main tab, click **System > Configuration**.
6. In the Geolocation area, click **Reload** in the **Operations** setting.
The system reloads the default geolocation database files that are stored in `/usr/share/GeoIP`.

Reloading default geolocation data using tmsh

To uninstall an update to the IP geolocation database, delete the RPM files, and then reload the default geolocation database files using `tmsh`.

1. At the BASH prompt, to query the RPM database and determine what geolocation data is installed, run this command:

```
rpm -qa --dbpath /shared/lib/rpm/
```

The system returns a list of RPMs, for example:

```
geoip-data-ISP-1.0.0-20110203.61.0  
geoip-data-Region2-1.0.0-20110203.61.0  
geoip-data-Org-1.0.0-20110203.61.0
```

2. To uninstall the RPMs, for each RPM in the list, run this command:

```
rpm -e --dbpath /shared/lib/rpm/ <name of file>
```

For example, to uninstall `geoip-data-ISP-1.0.0-20110203.61.0`, run this command: `rpm -e --dbpath /shared/lib/rpm/ geoip-data-ISP-1.0.0-20110203.61.0`

3. To remove the symlink in the `/shared/GeoIP` directory, run this command:

```
rm -f /shared/GeoIP/*
```

4. Log on to `tmsh`.

5. Run this command:

```
load / sys geoip
```

The system reloads the default geolocation database files that are stored in `/usr/share/GeoIP`.

Preparing BIG-IP DNS for Static Persist Load Balancing for Transaction-Oriented Traffic

About Static Persist load balancing

The Static Persist load balancing method uses the persist mask, with the source IP address of the LDNS, in a deterministic algorithm to send requests to a specific pool member (virtual server). Using this method, BIG-IP® DNS sends DNS name resolution requests to the first available pool member based on a hash algorithm that determines the order of the pool members. This algorithm orders the pool members differently for each LDNS that is sending requests to BIG-IP DNS, taking into account the Classless Inter-Domain Routing (CIDR) of the LDNS. As BIG-IP DNS distributes requests across all pool members, requests from each LDNS (and thus, each client) are generally sent to the same pool member. When the selected pool member becomes unavailable, BIG-IP DNS sends requests to another pool member. When the original pool member becomes available again, BIG-IP DNS sends requests to that pool member.

About configuring BIG-IP DNS to provide Static Persist load balancing

You can configure BIG-IP® DNS to provide Static Persist load balancing of LDNS to virtual servers. With this configuration, BIG-IP DNS guarantees that certain transactions are routed through a single transaction manager, such as Local Traffic Manager™, or another server array manager. This is beneficial for transaction-oriented traffic, such as e-commerce shopping carts, online trading, and online banking.

Configuring CIDR options for Static Persist load balancing

The Static Persist load balancing method uses the Classless Inter-Domain Routing (CIDR) system to identify which IP addresses require a persistent connection to a virtual server. You can configure the range of IP addresses by configuring the Static Persist CIDR setting for IPv4 and IPv6 addresses.

1. On the Main tab, click **DNS > Settings > GSLB > Load Balancing**.
The GSLB Load Balancing configuration screen opens.
2. For the **Static Persist CIDR (IPv4)** setting, type the number of bits that determine the IP address range for IPv4 addresses.
The default value is 32.
3. For the **Static Persist CIDR (IPv6)** setting, type the number of bits that determine the IP address range for IPv6 addresses.
The default value is 128.
4. Click **Update**.

You can now configure a pool for Static Persist load balancing.

Legal Notices

Legal notices

Publication Date

This document was published on June 28, 2018.

Publication Number

MAN-0593-01

Copyright

Copyright © 2018, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

For a current list of F5 trademarks and service marks, see <http://www.f5.com/about/guidelines-policies/trademarks/>.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at: <https://f5.com/about-us/policies/patents>

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Legal Notices

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index

B

BIG-IP DNS
 about 5
 introducing 5

C

CIDR
 configuring for Static Persist load balancing 35
 CNAME pools
 and Topology load balancing 31
 CNAME wide IPs
 and Topology load balancing 31
 Completion Rate load balancing method, about 14
 CPU load balancing method, about 14

D

database
 installing geolocation 31
 dependencies, and virtual server status 9
 Drop Packet load balancing method, about 13
 dynamic load balancing methods 7
 dynamic ratio load balancing 18
 Dynamic Ratio setting, about 18

F

Fallback IP load balancing method, about 15
 fallback load balancing method, and verifying virtual server availability 15

G

geolocation data
 about 31
 and topology records 31
 downloading and installing updates 31
 reloading default data 32
 reloading default data using tmsh 32
 Global Availability load balancing method, about 10
 global server load balancing, and virtual server dependency 9

H

Hops load balancing method, about 14

I

IP geolocation data
 and topology records 31
 downloading and installing updates 31

K

Kilobyte/Second load balancing method, about 13

L

Least Connections load balancing method, about 13
 limit settings, defined 10
 load balancing
 about pool-level 12
 about Topology 21
 about wide IPs and pool order 10
 and limit settings 10
 and resource availability 9
 and sort order of topology records 22
 and virtual server dependency 9
 using tiered 5
 using wide IP-level 5
 load balancing methods
 about Static Persist 35
 about Topology 21
 and dynamic ratio 18
 Completion Rate 14
 CPU 14
 customizing QoS 17
 Drop Packet 13
 dynamic 7
 Fallback IP 15
 Global Availability 10
 Hops 14
 Kilobyte/Second 13
 Least Connections 13
 None 15
 Packet Rate 13
 QoS 15
 Ratio 11
 Return to DNS 14
 Round Robin 12
 Round Trip Times 13
 static 5
 Static Persist 14, 35
 understanding QoS equation 15
 Virtual Server Capacity 13
 Virtual Server Score 13
 Longest Match sorting
 and topology records 22
 disabling 23
 enabling 24

M

manual resume feature
 and pools 18
 and virtual servers 19

N

None load balancing method, about 15

P

Packet Rate load balancing method, about 13

pool-level load balancing, about 12

pools

- and example Topology load balancing configuration 28

- and pool member weighting 11

- and the manual resume feature 18

- and Topology load balancing 30

- and Topology load balancing method 28

- organizing within wide IPs 10

- restoring availability manually 18

- weighting in wide IPs 11

preferred load balancing method, using when pool metrics are unavailable 18

Q

QoS equation

17

- about customizing 16

- understanding 15

QoS method

- about 15

- customizing equation 17

R

Ratio load balancing method

- about 11

- and pool member weighting 11

- and pool weighting in wide IPs 11

regions

- and Topology load balancing 21

- creating for Topology load balancing 22

resource availability, and load balancing 9

Return to DNS load balancing method, about 14

Round Robin load balancing method, about 12

Round Trip Times load balancing method, about 13

S

static load balancing methods 5

Static Persist load balancing

- and configuring CIDR options 35

Static Persist load balancing method

35

- about 14, 35

T

tiered load balancing 5

tmsh, and reloading default geolocation data 32

Topology load balancing

- about 21

- and an example for pools 28

- and an example for wide IPs 25

- and Longest Match sorting of records 22

- and pools 30

- and regions 21

- and wide IP addresses 25

- and wide IPs 27

Topology load balancing and creating regions 22

Topology load balancing method 21

topology records

- about 21

- and IP geolocation data 31

- and Longest Match sorting 22

- changing sort order 23

- configuring Longest Match sort order 24

- creating 24

- deleting 25

U

user-defined regions

- about 21

- creating for Topology load balancing 22

V

verifying virtual server availability, and fallback load balancing method 15

virtual server availability, verifying 15

Virtual Server Capacity load balancing method, about 13

virtual server dependency, and load balancing 9

virtual servers

- and configuring dependencies 9

- and weighting of pool members 11

- restoring availability manually 19

Virtual Server Score load balancing method, about 13

W

wide IP addresses

- and Topology load balancing 25

wide IP-level load balancing 5

wide IPs

- and example Topology load balancing configuration 25

- and pool order 10

- and pool weighting 11

- and Topology load balancing 30