

BIG-IP[®] DNS: Concepts

Version 12.0



Table of Contents

Legal Notices.....	7
Legal notices.....	7
About Global Server Load Balancing.....	9
Introducing BIG-IP DNS.....	9
About global server load balancing.....	9
Static load balancing methods.....	9
Dynamic load balancing methods.....	11
About load balancing and resource availability.....	13
About virtual server dependency.....	13
Configuring virtual server availability to be dependent on the status of other virtual servers.....	13
Limit settings for resource availability.....	14
About wide IP-level load balancing.....	14
About the Global Availability load balancing method.....	14
About the Ratio load balancing method.....	15
About the Round Robin load balancing method.....	16
About Topology load balancing.....	16
About pool-level load balancing.....	16
About the Drop Packet load balancing method.....	17
About the Virtual Server Score load balancing method.....	17
About the Virtual Server Capacity load balancing method.....	17
About the Round Trip Times load balancing method.....	17
About the Packet Rate load balancing method.....	17
About the Least Connections load balancing method.....	17
About the Kilobyte/Second load balancing method.....	17
About the Hops load balancing method.....	18
About the Completion Rate load balancing method.....	18
About the CPU load balancing method.....	18
About the Return to DNS load balancing method.....	18
About Static Persist load balancing.....	18
About the Fallback IP load balancing method.....	19
About the None load balancing method.....	19
About the QoS load balancing method.....	19
About dynamic ratio load balancing.....	22
Using the preferred load balancing method when metrics are unavailable.....	22
Configuring the resources in a pool for manual resume.....	22
Restoring availability of a pool member manually.....	23

Communications Between BIG-IP DNS and Other Systems.....	25
About establishing communications between BIG-IP DNS and other systems.....	25
About iQuery.....	25
About iQuery and communications between BIG-IP systems.....	25
Viewing iQuery statistics	26
About the gtm_add script.....	26
About the big3d_install script.....	27
About the bigip_add script.....	27
Configuration Synchronization.....	29
About configuration synchronization.....	29
About NTP servers and BIG-IP DNS configuration synchronization.....	29
BIG-IP DNS Configuration.....	31
About listeners.....	31
How listeners process network traffic.....	31
About wildcard listeners.....	32
About Prober pools.....	32
About Prober pool statistics.....	33
About Prober pool status.....	33
About probes.....	33
About BIG-IP DNS probes of an LDNS.....	33
Converting a statistics collection server to a Prober pool automatically.....	34
About delegation of LDNS probes.....	35
About LDNS entries on a BIG-IP DNS.....	35
Protocols and ports used by big3d during communications with local DNS servers.....	35
About wide IPs.....	36
About wildcard characters in wide IP names.....	36
About persistence connections.....	37
About wide IPs and a last resort pool.....	38
About data centers.....	39
About servers.....	39
About third-party host servers.....	39
About third-party load balancing servers.....	40
About virtual servers.....	40
About pools and pool members.....	41
About CNAME records.....	41
About links.....	41
Defining a link.....	42
About distributed applications.....	42
About ZoneRunner.....	43
About named.conf.....	43

Creating a master DNS zone.....	43
Creating a hint zone.....	44
Configuring BIG-IP DNS to allow zone file transfers.....	44
About DNS views.....	46
Types of DNS zone files.....	47
Types of DNS resource records.....	48
About DNSSEC.....	49
About DNSSEC keys.....	49
About enhancing DNSSEC key security.....	49
Viewing DNSSEC records in ZoneRunner.....	50
Protocols supported by the BIG-IP system.....	50

Legal Notices

Legal notices

Publication Date

This document was published on June 28, 2018.

Publication Number

MAN-0591-00

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.

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. 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 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.

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.

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. 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.

Communications Between BIG-IP DNS and Other Systems

About establishing communications between BIG-IP DNS and other systems

Before BIG-IP[®] DNS can operate as an integrated component within your network, communications must be established between BIG-IP DNS and the other BIG-IP systems with which BIG-IP DNS must exchange information. The three scripts described in the table are used to establish these communications in specific instances.

Script name	When to run
gtm_add	To install a new BIG-IP DNS on a network that includes a BIG-IP DNS synchronization group that contains at least one other BIG-IP DNS.
big3d_install	To install a new BIG-IP DNS on a network that includes BIG-IP systems running earlier versions of the BIG-IP system software.
bigip_add	To install a new BIG-IP DNS on a network that includes BIG-IP systems running the same version of the BIG-IP system software.

About iQuery

BIG-IP[®] systems use an XML protocol named *iQuery*[®] to communicate with other BIG-IP systems using gzip compression. BIG-IP systems must exchange SSL certificates and be members of the same configuration synchronization group before the systems can share information using iQuery.

Tip: *iqdump* is a command you can use to view the data transmitted between systems using iQuery.

Important: *BIG-IP systems send iQuery communications only on the VLAN on which the systems receive incoming messages.*

About iQuery and communications between BIG-IP systems

The `gtmd` agent on BIG-IP[®] DNS uses the *iQuery*[®] protocol to communicate with the local `big3d` agent, and the `big3d` agents installed on other BIG-IP systems. The `gtmd` agent monitors both the availability of the BIG-IP systems, and the integrity of the network paths between the systems that host a domain and the local DNS servers that attempt to connect to that domain.

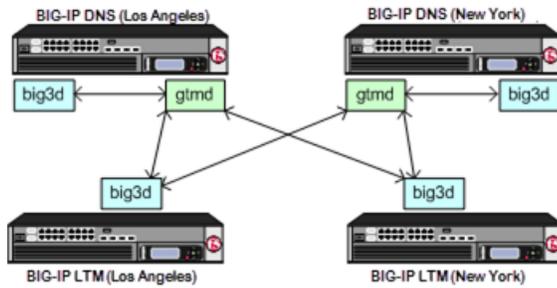


Figure 1: Communications between big3d and gtm agents using iQuery

Viewing iQuery statistics

Ensure that the BIG-IP[®] DNS configuration contains at least one BIG-IP server object with a self IP address.

To view information about the connections between BIG-IP DNS and other BIG-IP systems, view iQuery[®] statistics.

1. On the Main tab, click **Statistics > Module Statistics > DNS > GSLB**.
The Global Traffic statistics screen opens.
2. From the **Statistics Type** list, select **iQuery**.
Information about the iQuery connections between this system and other BIG-IP systems in your network displays.
3. When you want to estimate iQuery traffic throughput, click **Reset**.
The following statistics are reset to zero:
 - iQuery Reconnects
 - Bytes In
 - Bytes Out
 - Backlogs
 - Bytes Dropped

To view information about the iQuery[®] connections between a different BIG-IP DNS and the BIG-IP systems in your network, log in to that BIG-IP DNS and repeat this procedure.

About the gtm_add script

The `gtm_add` script is designed to be run on a BIG-IP[®] DNS you are installing on your network, when you want to integrate the system into a previously created BIG-IP DNS synchronization group that includes at least one other BIG-IP DNS. The script copies the remote BIG-IP DNS configuration to the local BIG-IP DNS system.

Note: Prior to running the `gtm_add` script, the BIG-IP DNS system you are adding to the network must be defined in the configuration of the BIG-IP DNS system in the configuration synchronization group to which you are adding the new system.

About the `big3d_install` script

The `big3d_install` script is designed to be run on a BIG-IP® DNS you are installing on a network that includes BIG-IP systems of earlier versions. The `big3d_install` script connects to each BIG-IP system on your network, extracts the IP addresses of the devices, and automatically updates the `big3d` agents on all devices.

About the `bigip_add` script

The `bigip_add` script is designed to run on a BIG-IP® DNS you are installing on a network that includes BIG-IP systems of the same version. The `bigip_add` script exchanges SSL certificates with each of the other BIG-IP systems to authorize communication between the devices.

Configuration Synchronization

About configuration synchronization

Configuration synchronization ensures the rapid distribution of BIG-IP® DNS settings to other BIG-IP systems that belong to the same BIG-IP DNS synchronization group. A *BIG-IP DNS synchronization group* might contain both BIG-IP DNS and BIG-IP® Link Controller™ systems.

Configuration synchronization occurs in the following manner:

- When a change is made to a BIG-IP DNS configuration, the system broadcasts the change to the other systems in the BIG-IP DNS synchronization group.
- When a configuration synchronization is in progress, the process must either complete or timeout, before another configuration synchronization can occur.

About NTP servers and BIG-IP DNS configuration synchronization

The Network Time Protocol (NTP) servers that BIG-IP DNS references ensure that each system in a BIG-IP DNS synchronization group is referencing the same time when verifying configuration file timestamps.

BIG-IP DNS Configuration

About listeners

A *listener* is a specialized virtual server that passively checks for DNS packets on port 53 and the IP address you assign to the listener. When a DNS query is sent to the IP address of the listener, BIG-IP DNS either handles the request locally or forwards the request to the appropriate resource.

How listeners process network traffic

You control how BIG-IP[®] DNS responds to DNS queries on a per-listener basis. The number of listeners you create depends on your network configuration and the destinations to which you want to send specific queries. For example, a single BIG-IP DNS can be the primary authoritative server for one domain, while forwarding other DNS queries to a different DNS server. BIG-IP DNS always manages and responds to DNS queries for the wide IPs that are configured on the system.

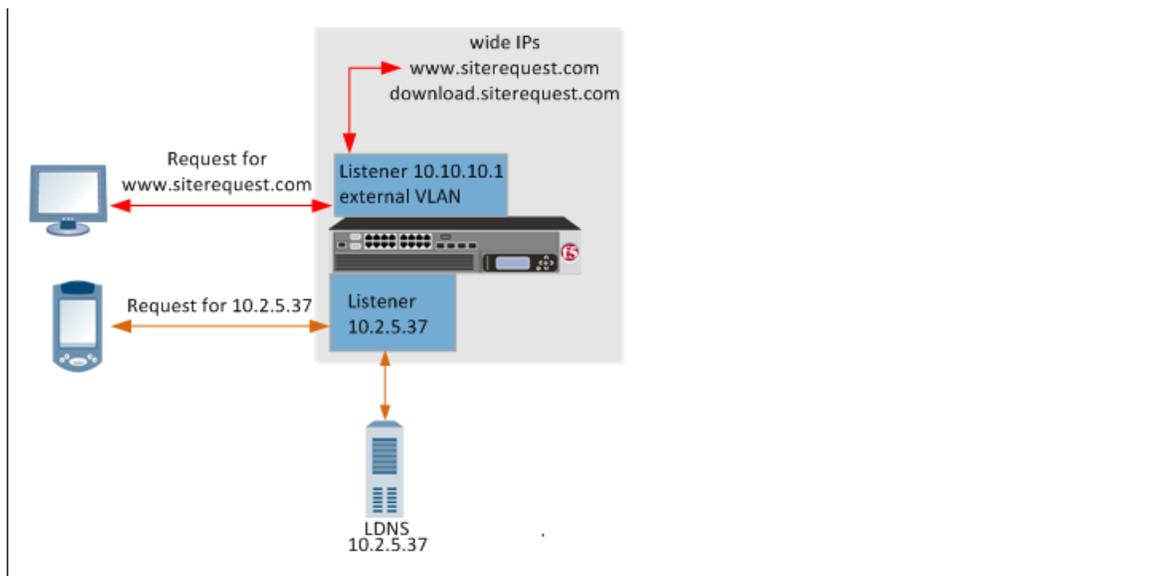
Consider the fictional company SiteRequest, once integrated into the network, BIG-IP DNS is responsible for:

- Managing and responding to requests for two wide IPs configured on the external VLAN:
 - `www.siterequest.com`
 - `downloads.siterequest.com`
- Forwarding DNS traffic destined for a DNS server at IP address `10.2.5.37`

In this scenario, BIG-IP DNS requires two listeners:

- A listener with an IP address that is on an external VLAN to manage DNS traffic destined for the wide IPs.
- A listener with the IP address of the local DNS server `10.2.5.37` to which the system can forward incoming traffic destined for that server.

Figure 2: Listener configuration example



About wildcard listeners

A *wildcard listener* is a special listener that is assigned an IP address of 0.0.0.0 and the DNS query port (port 53). When you want BIG-IP® DNS to respond to DNS queries coming into your network, regardless of the destination IP address of the given request, you use a wildcard listener.

About Prober pools

A *Prober pool* is an ordered collection of one or more BIG-IP® systems. BIG-IP DNS can be a member of more than one Prober pool, and a Prober pool can be assigned to an individual server or a data center. When you assign a Prober pool to a data center, by default, the servers in that data center inherit that Prober pool.

The members of a Prober pool perform monitor probes of servers to gather data about the health and performance of the resources on the servers. BIG-IP DNS makes load balancing decisions based on the gathered data. If all of the members of a Prober pool are marked down, or if a server has no Prober pool assigned, BIG-IP DNS reverts to a default intelligent probing algorithm to gather data about the resources on the server.

This figure illustrates how Prober pools work. BIG-IP DNS contains two BIG-IP Local Traffic Manager™ (LTM™) systems that are assigned Prober pools and one BIG-IP LTM system that is not assigned a Prober pool:

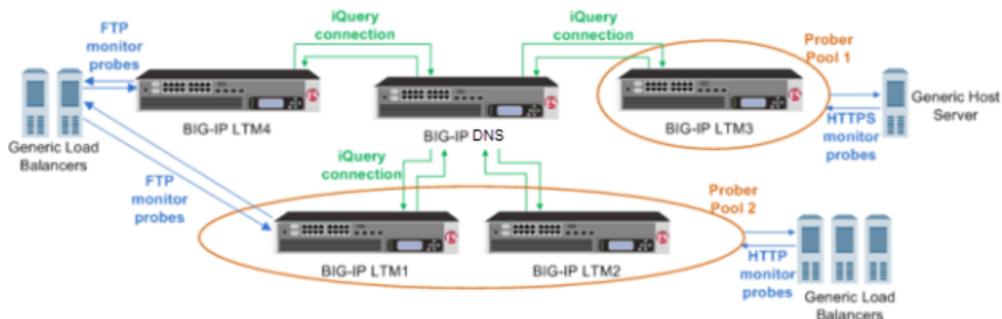


Figure 3: BIG-IP systems with prober pools

Prober Pool 1 is assigned to a generic host server

BIG-IP LTM3 is the only member of Prober Pool 1, and performs all HTTPS monitor probes of the server.

Prober Pool 2 is assigned to generic load balancers

BIG-IP LTM1 and BIG-IP LTM2 are members of Prober Pool 2. These two systems perform HTTP monitor probes of generic load balancers based on the load balancing method assigned to Prober Pool 2.

The generic load balancers on the left side of the graphic are not assigned a Prober pool

BIG-IP DNS can solicit any BIG-IP system to perform FTP monitor probes of these load balancers, including systems that are Prober pool members.

About Prober pool statistics

You can view the number of successful and failed probe requests that the BIG-IP® DNS system (on which you are viewing statistics) made to the Prober pools. These statistics reflect only the number of Probe requests and their success or failure. These statistics do not reflect the actual probes that the pool members made to servers on your network.

Prober pool statistics are not aggregated among the BIG-IP DNS systems in a synchronization group. The statistics on one BIG-IP DNS include only the requests made from that BIG-IP DNS system.

About Prober pool status

The status of a Prober pool also indicates the status of the members of the pool. If at least one member of a Prober pool has green status (Available), the Prober pool has green status.

The status of a Prober pool member indicates whether the BIG-IP DNS system, on which you are viewing status, can establish an iQuery connection with the member.

***Note:** If a Prober pool member has red status (Offline), no iQuery connection exists between the member and the BIG-IP DNS system on which you are viewing status. Therefore, that BIG-IP DNS system cannot request that member to perform probes, and the Prober pool will not select the member for load balancing.*

About probes

A *probe* is an action a BIG-IP® system takes to acquire data from other network resources. BIG-IP DNS uses probes to track the health and availability of network resources.

About BIG-IP DNS probes of an LDNS

BIG-IP® DNS is responsible for acquiring data from local DNS servers (LDNS) using probes. Unlike probes conducted on internal systems, such as web servers, probes of an LDNS require that BIG-IP DNS verifies data from a resource that exists outside the network. Typically, this data is the path information BIG-IP DNS requires when conducting Quality of Service, Round Trip Time, Completion Rate, and Hops load balancing methods.

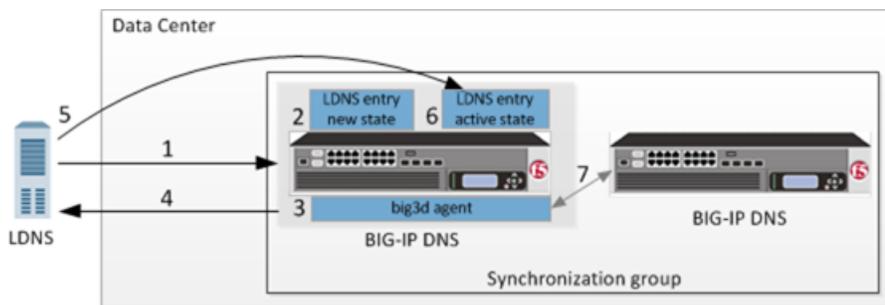


Figure 4: Flow of probe process

The steps in the probe process

1. An LDNS sends a DNS name resolution request to BIG-IP DNS for a wide IP.
2. The specific BIG-IP DNS in a synchronization group that is chosen to manage probing of a resource creates an entry with a state of New (unless an entry for the LDNS already exists).

Note: Once a specific BIG-IP DNS system becomes responsible for managing the probing of a resource, that BIG-IP DNS maintains that responsibility until it goes offline, a new BIG-IP DNS is added to the data center, or the configuration of the resource changes.

3. The BIG-IP DNS that is responsible for managing the probing of the resource, delegates the probe to a BIG-IP DNS that is currently responsible for conducting the fewest number of probes.

Note: BIG-IP DNS checks this statistic each time a probe of a specific resource is required; as a result, the BIG-IP DNS that probes the resource can change from one probe instance another.

4. The `big3d` agent on the BIG-IP DNS that is delegated to probe the resource, sends a probe to the LDNS.
5. The LDNS responds to the probe.
6. BIG-IP DNS updates the LDNS entry, assigning it an Active state.
7. The `big3d` agent then broadcasts the results of the probe to all BIG-IP DNS systems in the synchronization group.

Converting a statistics collection server to a Prober pool automatically

In version 10.2 of BIG-IP® DNS, you could assign a single BIG-IP® system to probe a server to gather health and performance data. You did this by specifying the IP address of the BIG-IP system (which you chose to perform probes of the server) in the **Statistics Collection Server** field of the server. In version 11.0, this feature was replaced by the Prober pool feature.

When you upgrade from version 10.2.x to version 11.x, if a single BIG-IP system was assigned to probe a server, BIG-IP DNS converts the single server to a Prober pool with one member, and then assigns the Prober pool to the server to which the Statistics Collection server was originally assigned. The name of the new Prober pool is based on the IP address of the original Statistics Collection server. If the original Statistics Collection server had an IP address of 10.10.2.3, the name of the automatically created Prober pool is `prober_pool_10_10_2_3`.

About delegation of LDNS probes

By default, BIG-IP DNS delegates the probe of a resource to a BIG-IP DNS that is in the same data center as the resource, because the close proximity improves probe response time. However, when more than one BIG-IP DNS in a synchronization group resides in the same data center, an algorithm is used to delegate the probes of resources equally among those BIG-IP DNS systems. This ensures that no BIG-IP DNS becomes overloaded with conducting probes, which can cause a decrease in the performance of the other tasks for which the BIG-IP system is responsible. When a data center does not contain a BIG-IP DNS, an algorithm is used to distribute the actual probes of a resource in that data center equally among the BIG-IP DNS systems in the synchronization group.

About LDNS entries on a BIG-IP DNS

An LDNS entry can be in one of three states: New, Pending, or Active. In general, the New and Pending states are temporary. An LDNS entry remains in one of these states only until the LDNS responds to the first probe request from a BIG-IP DNS.

New

This LDNS has not contacted BIG-IP DNS before.

Pending

This LDNS has contacted BIG-IP DNS before; however, this LDNS has yet to respond to a probe.

Active

BIG-IP DNS has an entry for this LDNS.

Protocols and ports used by big3d during communications with local DNS servers

This table describes the protocols and ports the `big3d` agent uses to communicate with an LDNS when collecting path data for the local DNS servers.

Table 2: Communication between big3d agents and local DNS servers

From	To	Protocol	From port	To port	Purpose
big3d agent	LDNS	ICMP	n/a	n/a	Probe using ICMP pings
big3d agent	LDNS	TCP	>1023	53	Probe using TCP (Cisco® routers: allow establish)
LDNS	big3d agent	TCP	53	1023	Replies using TCP (Cisco® routers: allow establish)
big3d agent	LDNS	UDP	53	33434	Probe using UDP or the traceroute utility
LDNS	big3d agent	ICMP	n/a	n/a	Replies to ICMP, UDP

From	To	Protocol	From port	To port	Purpose
					pings, or traceroute utility probes
big3d agent	LDNS	dns_rev, dns_dot	>1023	53	Probe using DNS rev or DNS dot
big3d agent	LDNS	dns_rev, dns_dot	53	>1023	Replies to DNS rev or DNS dot probes

About wide IPs

A *wide IP* maps a fully-qualified domain name (FQDN) to one or more pools of virtual servers that host the content of a domain. When an LDNS issues a DNS name resolution for a wide IP, the configuration of the wide IP indicates which pools of virtual servers are eligible to respond to the request, and which load balancing methods BIG-IP DNS uses to select the pool.

About wildcard characters in wide IP names

BIG-IP[®] DNS supports these wildcard characters in wide IP names and aliases:

Question mark (?)

Use to replace a single character, except a dot (.).

Asterisk (*)

Use to replace multiple consecutive characters with the exception of dots (.)

You can use one or more question marks or asterisks, or both question marks and asterisks in a wide IP name or alias.

Valid uses of wildcard characters for the wide IP name `www.mydomain.net` include:

- `???.mydomain.net`
- `www.??domain.net`
- `www.my*.net`
- `www.??*.net`
- `www.my*.*`
- `???.my*.*`
- `*.*.net`
- `www.*.??`

Using wildcard characters in wide IPs to minimize maintenance tasks

Determine the domain names and aliases for which you want to configure wide IPs.

Create a wide IP using wildcard characters in the name to represent a domain when you have a large quantity of aliases that you want to use for the domain.

1. On the Main tab, click **DNS > GSLB > Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.

***Tip:** You can use two different wildcard characters in the wide IP name: asterisk (*) to represent several characters and question mark (?) to represent a single character. This reduces the number of aliases you have to add to the configuration.*

4. From the **Pool** list, select the pools that this wide IP uses for load balancing.
The system evaluates the pools based on the wide IP load balancing method configured.
 - a) From the **Pool** list, select a pool.
A pool can belong to more than one wide IP.
 - b) Click **Add**.
5. Click **Finished**.

About persistence connections

Most load balancing methods divide DNS name resolution requests among available pools or virtual servers. Each time BIG-IP DNS receives a request, it sends that request to the most appropriate resource based on the configuration of your network.

For example, when a user visits a web site, multiple DNS name resolution requests are generated as the user moves from page to page. Depending on the load balancing method configured, BIG-IP DNS sends each request to a different server, virtual server, or data center. In certain circumstances, you might want to ensure that a user remains with a given set of resources throughout the session. For example, a user attempting to conduct an online banking transaction needs to remain with the same set of resources to ensure that the transaction is completed successfully.

Configuring BIG-IP DNS for persistent connections

Configure BIG-IP DNS for persistent connections when you want a user to stay with a specific set of resources during a web transaction.

1. On the Main tab, click **DNS > GSLB > Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.

***Tip:** You can use two different wildcard characters in the wide IP name: asterisk (*) to represent several characters and question mark (?) to represent a single character. This reduces the number of aliases you have to add to the configuration.*

4. From the **Load Balancing Method** list, select the load balancing method the wide IP uses to select a pool.
5. On the menu bar, click **Pools**.
6. From the **Persistence** list, select **Enabled**.

7. In the **Persistent TTL** field, type the number of seconds the persistence entry is valid.
This value can range from 0 to 4294967295 seconds.
8. Click **Finished**.

Once a resource has sent a response to a DNS name resolution request, BIG-IP DNS sends subsequent requests from the same connection to that same resource until the current session is completed.

About wide IPs and a last resort pool

BIG-IP[®] DNS considers any pool associated with a wide IP as a potential resource to which to load balance DNS name resolution requests. A *last resort pool* is a pool of virtual servers to which BIG-IP DNS sends DNS name resolution requests in the event that all other pools associated with the wide IP are unavailable. When you design your network, ensure that this particular pool is kept in reserve and not used as part of the normal global server load balancing.

Configuring a wide IP with a last resort pool

Ensure that a pool of virtual servers that is not part of the normal global server load balancing operations exists in the configuration.

Create a wide IP to map a FQDN to one or more pools of virtual servers that host the content of the domain, and assign a last resort pool to the wide IP.

1. On the Main tab, click **DNS > GSLB > Wide IPs**.
The Wide IP List screen opens.
2. Click **Create**.
The New Wide IP screen opens.
3. In the **Name** field, type a name for the wide IP.

***Tip:** You can use two different wildcard characters in the wide IP name: asterisk (*) to represent several characters and question mark (?) to represent a single character. This reduces the number of aliases you have to add to the configuration.*

4. From the **Pool** list, select the pools that this wide IP uses for load balancing.
The system evaluates the pools based on the wide IP load balancing method configured.
 - a) From the **Pool** list, select a pool.
A pool can belong to more than one wide IP.
 - b) Click **Add**.
5. From the **Last Resort Pool** list, select a pool for the system to use when no other resources are available.
6. Click **Finished**.

BIG-IP DNS uses the resources included in the last resort pool only if no other resources are available to handle DNS name resolution requests for the wide IP.

About data centers

All of the resources on your network are associated with a data center. BIG-IP® DNS consolidates the paths and metrics data collected from the servers, virtual servers, and links in the data center. BIG-IP DNS uses that data to conduct load balancing and route client requests to the best-performing resource based on different factors.

BIG-IP DNS might send all requests to one data center when another data center is down. Alternatively, BIG-IP DNS might send a request to the data center that has the fastest response time. A third option might be for BIG-IP DNS to send a request to the data center that is located closest to the client's source address.

Tip: The resources associated with a data center are available only when the data center is also available.

About servers

A *server* defines a physical system on the network. Servers contain the virtual servers that are the ultimate destinations of DNS name resolution requests. BIG-IP® DNS supports three types of server,s as shown in the table.

BIG-IP® systems

Any member of the BIG-IP system product line.

Third-party load balancing systems

A third-party load balancing system is any system, other than a BIG-IP system, that supports and manages virtual servers on the network.

Third-party host servers

A third-party host server is a resource to which the BIG-IP system load balances DNS traffic, for example, a web server, file server, or SQL server.

About third-party host servers

A *host* is a network resource that is not a part of the BIG-IP® product family and does not provide load balancing. BIG-IP® DNS supports these host servers:

- CacheFlow®
- NetApp™
- Sun Solaris™
- Windows 2000 Server (You can monitor the Windows Vista® Enterprise Server using the Windows 2000 Server.)
- Windows NT 4.0™

About third-party load balancing servers

BIG-IP[®] DNS interacts with other load balancing servers to determine availability and assess performance when responding to DNS name resolution requests. BIG-IP DNS supports these load balancing servers:

- Alteon[®] Ace Director
- Cisco[®] CSS
- Cisco[®] LocalDirector v2
- Cisco[®] LocalDirector v3
- Cisco[®] SLB
- Extreme
- Foundry[®] ServerIron
- Radware WSD
- Other generic load balancers

About virtual servers

A *virtual server* is a specific IP address and port number that points to a resource on the network. In the case of host servers, this IP address and port number likely point to the resource itself. With load balancing systems, virtual servers are often proxies that allow the load balancing server to manage a resource request across a multitude of resources.

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.

Configuring virtual server status for clusters

You can configure virtual server status to be dependent only on the timeout value of the monitor associated with the virtual server. This ensures that when the primary blade in a cluster becomes unavailable, the `gtmd` agent on the new primary blade has time to establish new iQuery® connections with and receive updated status from other BIG-IP® systems.

Tip: The `big3d` agent on the new primary blade must be up and functioning within 90 seconds (the timeout value of the BIG-IP monitor).

1. On the Main tab, click **DNS > Settings > GSLB > General**.
The General configuration screen opens.
2. Select **Depends on Monitors Only** from the **Virtual Server Status** list.
3. Click **Update**.

About pools and pool members

A *pool* is a collection of virtual servers that can reside on multiple servers. A *virtual server* is a combination of IP address and port number that points to a specific resource on the network. When you add a virtual server to a pool, it becomes a pool member. A *pool member* is a virtual server that has attributes that pertain to the virtual server only in the context of the pool. A virtual server can be a member of multiple pools and have different attributes in each pool. BIG-IP® DNS directs traffic to a pool member, based on the attributes of the pool member.

To illustrate the difference between pool members and virtual servers, consider the fictional company SiteRequest. In the London data center, the IT team has a virtual server that acts as a proxy for a BIG-IP Local Traffic Manager™ (LTM™). This virtual server is the primary resource for DNS name resolution requests for the company web page that originate from Europe. This virtual server is also the backup resource for requests that originate from the United States. Because these are two distinctly different roles, the virtual server is a pool member in two different pools. The IT team can use this configuration to customize the virtual server for each pool to which it belongs, without modifying the actual virtual server itself.

About CNAME records

A *CNAME* record specifies that a domain name is an alias of another domain. When you create a CNAME pool with members specifying the canonical name, BIG-IP® DNS responds to DNS name resolution requests for A or CNAME name record type queries with these names if this pool is added to the said wide IP.

About links

A *link* is a logical representation of a physical device (router) that connects your network to the Internet. BIG-IP® DNS tracks the performance of links, which influence the availability of pools, data centers, wide IPs, and distributed applications.

Defining a link

Ensure that at least one data center exists in the configuration.

Gather information about the routers that you want to define as links, including:

- IP addresses
- Data center location

Define links to aid BIG-IP® DNS in determining resource availability.

1. On the Main tab, click **DNS > GSLB > Links**.

The Links List screen opens.

2. Click **Create**.

The New Link screen opens.

3. Type a name for the link.

Important: Link names are limited to 63 characters.

4. Specify whether the link uses address translation when communicating between the network and the Internet.

Important: If you enable this setting, the BIG-IP link monitor cannot monitor outbound traffic through this link.

5. Type the IP address of a router in the **Address** field, and then click **Add**.

You can add more than one IP address, depending on how the server on which you are creating the link interacts with the rest of your network.

6. Select the data center where the router that the link represents resides.

7. In the **Uplink Address** field, specify the IP address of the router on the ISP side of the link.

When you configure an uplink address, the BIG-IP system sends SNMP requests to the IP addresses configured in the **Router Address List**. The system uses the statistics that the router returns to distinguish between internal-only traffic and traffic destined for the Internet.

8. Assign the BIG-IP Link monitor to the link by moving it from the **Available** list to the **Selected** list.

9. Click **Create**.

The `big3d` agent can now gather and analyze path and metrics information about outbound traffic passing through the router the link represents.

About distributed applications

A *distributed application* is a collection of one or more wide IPs, data centers, and links that serve as a single application to a web site visitor. Configuring a distributed application provides several advantages:

- You can organize logical network components into groups that represent a business environment.
- You can configure a distributed application to be dependent upon the availability of a data center, server, or link. This dependency ensures that a user cannot access a distributed application when a portion of the resources are unavailable.
- You can define persistence for the distributed application, ensuring that a user, who accesses the distributed application uses the same resources during a single session.

If the New York data center goes offline, a wide IP in that data center becomes unavailable. A distributed application associated with that wide IP also becomes unavailable. Consequently, the system does not send resolution requests to any of the distributed application resources, until the entire application becomes available again.

About ZoneRunner

You can use the ZoneRunner™ utility to create and manage DNS zone files and configure the BIND instance on BIG-IP® DNS (formerly GTM). With the ZoneRunner utility, you can:

- Import and transfer DNS zone files
- Manage zone resource records
- Manage views
- Manage a local nameserver and the associated configuration file, `named.conf`
- Transfer zone files to a nameserver
- Import only primary zone files from a nameserver

About named.conf

`named.conf` contains the primary operational characteristics of BIND, including DNS views, access control list definitions, and zones. The ZoneRunner™ utility updates `named.conf` when you modify the local BIND instance.

Using ZoneRunner to configure named.conf

Ensure that at least one zone is configured on BIG-IP® DNS.

Use ZoneRunner™ to edit `named.conf`, to decrease the risk of a syntax error that prevents the BIND system from performing as expected. Zonerunner provides an automatic syntax check and displays error messages to help you write the correct syntax.

1. On the Main tab, click **DNS > Zones > ZoneRunner > named Configuration**.
The named Configuration screen opens.
2. In the Options area, type additional configurations per your network design.
3. Click **Update**.

Creating a master DNS zone

A master zone is authoritative. Create a zone when you want to use ZoneRunner™ to manage DNS zones and resource records.

Tip: The BIG-IP® system can be either a primary or secondary DNS server.

1. On the Main tab, click **DNS > Zones > ZoneRunner > Zone List**.
The Zone List screen opens.

2. Click **Create**.
The New Zone screen opens.
3. From the **View Name** list, select **external**.
The external view is a default view to which you can assign zones.
4. In the **Zone Name** field, type a period character (.).
5. From the **Zone Type** list, select **Master**.
6. Clear the **Zone File Name** field, and type the zone file name.
db.external.siterequest.com

Note: Do not include a trailing dot.

7. In the Records Creation area, type the values for the SOA and NS record parameters.
8. Click **Finished**.

If you want further help creating a custom zone file, see *SOL8380* on www.askf5.com for instructions.

Creating a hint zone

Hint zones designate a subset of the root nameservers list. When the local nameserver starts (or restarts), the nameserver queries the root servers in the hint zone for the most current list of root servers. The root hint is built into BIND version 9.0 and later.

Create a zone when you want to use ZoneRunner™ to manage DNS zones and resource records.

Tip: The BIG-IP® system can be either a primary or secondary DNS server.

1. On the Main tab, click **DNS > Zones > ZoneRunner > Zone List**.
The Zone List screen opens.
2. Click **Create**.
The New Zone screen opens.
3. From the **View Name** list, select **external**.
The external view is a default view to which you can assign zones.
4. In the **Zone Name** field, type a period character (.).
5. From the **Zone Type** list, select **Hint**.
6. Clear the **Zone File Name** field, and type the zone file name.
db.external.siterequest.com

Note: Do not include a trailing dot.

7. Click **Finished**.

If you want further help creating a custom hint file, see *SOL8380* on www.askf5.com for instructions.

Configuring BIG-IP DNS to allow zone file transfers

By default, BIG-IP® DNS is configured to secure BIND to not allow zone transfers except from the localhost. However, you can configure BIG-IP DNS to allow zone file transfers to other DNS servers.

1. On the Main tab, click **DNS > Zones > ZoneRunner > named Configuration**.
The named Configuration screen opens.
2. In the **Options** field, modify the allow-transfer statement to include the IP address of the BIG-IP DNS.
You can modify the following allow-transfer statement to use the IP address of the BIG-IP DNS.

```
allow-transfer {
    localhost;
    192.168.10.105;
};
```

3. On the menu bar, click **View List**.
The View List screen opens.
4. Click the name of the view that contains the zone you are configuring.
The View Configuration screen opens.
5. In the Options area, modify the match-clients statement based on your configuration.

View configuration type

Single view configuration

Add to match-clients statement

```
view "external" {
    match-clients {
        "zrd-acl-000-000";
    };
};
```

Multiple view configuration, where you want to allow transfers from BIG-IP DNS

Modify the following match-clients statement to use the IP address of the BIG-IP DNS.

```
acl "internal-acl"
{ <IP address> ;
};

view "internal" {
    match-clients {
        "zrd-acl-000-001";
        "internal-acl";
        <IP address> ;
    };
};

view "external" {
    match-clients {
        "zrd-acl-000-000";
    };
};
```

6. Click **Update**.

To verify that zone transfers are working properly, modify this Linux command and run it on an external computer: `dig @<IP address> es.net. axfr`

The command should return a response similar to this:

```
; <<>> DiG? 9.5.0-P2 <<>> @192.17.1.253 es.net. axfr
; (1 server found)

;; global options: printcmd
```

```
es.net. 500 IN SOA siterequest.com.  
hostmaster.siterequest.com. 6 10800 3600 604800 60  
es.net. 500 IN NS siterequest.com.  
a.es.net. 30 IN A 192.17.1.100  
b.es.net. 30 IN A 192.18.1.100  
es.net. 500 IN SOA siterequest.com.  
hostmaster.siterequest.com. 6 10800 3600 604800 60  
;; Query time: 6 msec  
;; SERVER: 192.17.1.253#53(192.17.1.253)  
;; WHEN: Fri Mar 11 17:20:25 2011  
;; XFR size: 5 records (messages 1, bytes 180)
```

About DNS views

A DNS *view* is a modification of a nameserver configuration based on the community attempting to access it. Using views, you can build multiple nameserver configurations on the same server, and have those configurations apply dynamically when the request originates from a specified source.

If your DNS handles requests from both inside and outside your company, you can create two views: internal and external.

Creating a DNS view

It is helpful to keep in mind that ZoneRunner™ contains a default view named: external.

Create an additional DNS view to modify the local nameserver configuration to allow a specific community to access it.

1. On the Main tab, click **DNS > Zones > ZoneRunner > View List**.
The View List screen opens.
2. Click **Create**.
3. In the **View Name** field, type a name for the view.
4. From the **View Order** list, make a selection.

Option	Description
First	In the view hierarchy, this view is listed first.
Last	In the view hierarchy, this view is listed last.
After	In the view hierarchy, this view is listed immediately following the view that you select from the View List.

5. In the Options area, modify the match-clients statement based on your configuration.

View configuration type**Single view configuration****Add to match-clients statement**

```
view "external" {
    match-clients {
        "zrd-acl-000-000";
        any;
    };
};
```

Multiple view configuration, where you want to allow transfers from BIG-IP DNS

Modify the following match-clients statement to use the IP address of the BIG-IP DNS.

```
acl "internal-acl"
{ <IP address> ;
};

view "internal" {
    match-clients {
        "zrd-acl-000-001";
        "internal-acl";
        <IP address> ;
    };
};

view "external" {
    match-clients {
        "zrd-acl-000-000";
        any;
    };
};
```

6. In the Options area, type additional configurations per your network design.
7. Click **Finished**.

Types of DNS zone files

This table describes the types of DNS zone files.

DNS file type	Description
Primary	Zone files for a primary zone contain, at minimum, the start of authority (SOA) and nameserver (NS) resource records for the zone. Primary zones are authoritative, that is, they respond to DNS queries for the domain or sub-domain. A zone can have only one SOA record, and must have at least one NS record.
Secondary	Zone files for a secondary zone are copies of the principal zone files. At an interval specified in the SOA record, secondary zones query the primary zone to check for and obtain updated zone data. A secondary zone responds authoritatively for the zone provided that the zone data is valid.
Stub	Stub zones are similar to secondary zones, except that stub zones contain only the NS records for the zone. Note that stub zones are a specific feature of the BIND implementation of DNS. F5 Networks recommends that you use stub zones only if you have a specific requirement for this functionality.
Forward	The zone file for a forwarding zone contains only information to forward DNS queries to another nameserver on a per-zone (or per-domain) basis.
Hint	The zone file for a hint zone specifies an initial set of root nameservers for the zone. Whenever the local nameserver starts, it queries a root nameserver in the hint zone file to obtain the most recent list of root nameservers. Zone file import.

Types of DNS resource records

This table describes the types of DNS resource records that ZoneRunner™ supports.

DNS file type	Description
SOA (Start of authority)	The start of authority resource record, SOA, starts every zone file and indicates that a nameserver is the best source of information for a particular zone. The SOA record indicates that a nameserver is authoritative for a zone. There must be exactly one SOA record per zone. Unlike other resource records, you create a SOA record only when you create a new master zone file.
A (Address)	The Address record, or A record, lists the IP address for a given host name. The name field is the host's name, and the address is the network interface address. There should be one A record for each IP address of the machine.
AAAA (IPv6 Address)	The IPv6 Address record, or AAAA record, lists the 128-bit IPv6 address for a given host name.
CNAME (Canonical Name)	The Canonical Name resource record, CNAME, specifies an alias or nickname for the official, or canonical, host name. This record must be the only one associated with the alias name. It is usually easier to supply one A record for a given address and use CNAME records to define alias host names for that address.
DNAME (Delegation of Reverse Name)	The Delegation of Reverse Name resource record, DNAME, specifies the reverse lookup of an IPv6 address. These records substitute the suffix of one domain name with another. The DNAME record instructs DNS (BIG-IP® DNS, formerly GTM) (or any DNS server) to build an alias that substitutes a portion of the requested IP address with the data stored in the DNAME record.
HINFO (Host Information)	The Host Information resource record, HINFO, contains information on the hardware and operating system relevant to BIG-IP DNS (formerly GTM) (or other DNS).
MX (Mail Exchanger)	The Mail Exchange resource record, MX, defines the mail system(s) for a given domain.
NAPTR (Name Authority Pointer)	The Name Authority Pointer record, NAPTR, aids in the standardization of Uniform Resource Names (URNs). NAPTR records map between sets of URNs, URLs and plain domain names and suggest to clients the protocols available for communication with the mapped resource.
NS (nameserver)	The nameserver resource record, NS, defines the nameservers for a given domain, creating a delegation point and a subzone. The first name field specifies the zone that is served by the nameserver that is specified in the nameservers name field. Every zone needs at least one nameserver.
PTR (Pointer)	A name pointer resource record, PTR, associates a host name with a given IP address. These records are used for reverse name lookups.
SRV (Service)	The Service resource record, SRV, is a pointer with which an alias for a given service is redirected to another domain. For example, if the fictional company Site Request has an FTP archive hosted on archive.siterequest.com, the IT department can create an SRV record with which the alias ftp.siterequest.com is redirected to archive.siterequest.com.
TXT (Text)	The Text resource record, TXT, allows you to supply any string of information, such as the location of a server or any other relevant information that you want available.

About DNSSEC

Domain Name System Security Extensions (DNSSEC) is an industry-standard protocol that functions as an extension to the Domain Name System (DNS) protocol. BIG-IP® DNS uses DNSSEC to guarantee the authenticity of DNS responses, including zone transfers, and to return Denial of Existence responses thus protecting your network against DNS protocol and DNS server attacks.

About DNSSEC keys

BIG-IP® DNS, formerly Global Traffic Manager™ (GTM™), uses two types of DNSSEC keys to return DNSSEC-compliant responses: a *zone-signing key* to sign all of the records in a DNSSEC resource record set, and a *key-signing key* to sign only the DNSKEY record (that is the zone-signing key) of a DNSSEC record set.

About enhancing DNSSEC key security

To enhance DNSSEC key security, when automatic key management is configured, BIG-IP® DNS uses an automatic key rollover process that uses overlapping generations of a key to ensure that BIG-IP DNS can always respond to queries with DNSSEC-compliant responses. BIG-IP DNS dynamically creates new generations of each key based on the values of the **Rollover Period** and **Expiration Period** of the key.

The first generation of a key has an ID of 0 (zero). Each time BIG-IP DNS dynamically creates a new generation of a key, the ID increments by one. Over time, each generation of a key overlaps the previous generation of the key ensuring that BIG-IP DNS can respond to a DNSSEC query even if one generation of a key becomes unavailable. When a generation of a key expires, BIG-IP DNS automatically removes that generation of the key from the configuration. The value of the **TTL (time-to-live)** of a key specifies how long a client resolver can cache the key.

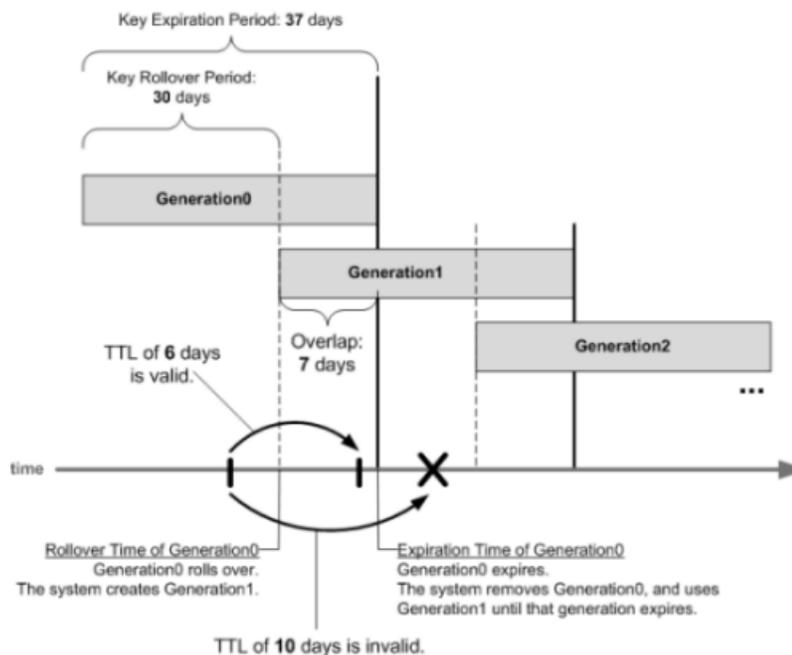


Figure 5: Overlapping generations of a key

Viewing DNSSEC records in ZoneRunner

Ensure that all DNSSEC records are added to the BIND configuration.

View the DNSSEC records using ZoneRunner™ when you want to evaluate how your network is handling DNSSEC traffic.

1. On the Main tab, click **DNS > Zones > ZoneRunner > Resource Record List**.
The Resource Record List screen opens.
2. From the **View Name** list, select the name of the view that contains the resource records you want to view.
3. From the **Zone Name** list, select the zone for which you want to view resource records.
4. From the **Type** list, select the type of resource records you want to view.
5. Click **Search**.

View the resource records that display.

Protocols supported by the BIG-IP system

This table describes the protocols supported by the BIG-IP® system.

Protocol	Description
iQuery® protocol	The <code>gtmd</code> agent on BIG-IP DNS, formerly Global Traffic Manager™ (GTM™), uses iQuery® to communicate with the local <code>big3d</code> agent and the <code>big3d</code> agents installed on other BIG-IP systems.
DNS	The BIG-IP system supports the Domain Name System (DNS) for distribution of DNS name resolution requests from clients and their local DNS servers to resources on your global network.
DNSSEC	BIG-IP DNS supports the DNS Security Extensions for secure zone signing and authentication of DNS responses.
HTTPS	BIG-IP supports Hypertext Transfer Protocol Secure (HTTPS) for secure web browsing.
SSL	The web server, which hosts the web-based Configuration utility, supports Secure Sockets Layer (SSL) connections as well as user authentication.
SNMP	BIG-IP supports the Simple Network Management Protocol (SNMP) for monitoring network resources.
SMTP	BIG-IP supports the Simple Mail Transfer Protocol (SMTP) for email transmission across the Internet.
SSH	BIG-IP supports Secure Shell (SSH) administrative connections for remote administration from the command line.
TCP wrappers	BIG-IP supports the use of TCP wrappers to provide an extra layer of security for network connections.
FTP	BIG-IP supports the File Transfer Protocol (FTP) for secure access to BIG-IP system software downloads from a web server.

Index

B

- big3d agent
 - and communication with an LDNS 35
 - and iQuery 25
 - and monitor timeout values 41
 - communicating with an LDNS 33
- BIG-IP DNS
 - about 9
 - about establishing inter-device communication 25
 - introducing 9
- BIG-IP systems, and iQuery connections 26

C

- canonical names
 - and pools 41
- clusters, configuring 41
- CNAME records
 - about 41
- communications, about establishing inter-device 25
- Completion Rate load balancing method, about 18
- configuration synchronization, about 29
- connections
 - viewing iQuery statistics 26
 - viewing status 26
- CPU load balancing method, about 18

D

- data acquisition
 - and BIG-IP system probe responsibility 35
 - and iQuery 25
 - and probes 33
- data centers
 - about 39
- dependencies, and virtual server status 13, 40
- distributed applications, defined 42
- DNSSEC, about 49
- DNSSEC keys, about 49
- DNSSEC records, viewing 36, 50
- DNS servers, and zone transfers 44
- DNS views, creating 46
- DNS zone files, described 47–48
- Drop Packet load balancing method, about 17
- dynamic load balancing methods 11
- dynamic ratio load balancing 22
- Dynamic Ratio setting, about 22

F

- Fallback IP load balancing method, about 19
- fallback load balancing method, and verifying virtual server availability 19

G

- Global Availability load balancing method, about 14
- global server load balancing, and virtual server dependency 13
- global traffic management
 - and listeners 31
 - and wildcard listeners 32
- gtm_add utility, about 26–27
- gtmd agent, and iQuery 25

H

- Hint zone, configuring using ZoneRunner 44
- Hops load balancing method, about 18
- hosts, defined 39

I

- iQuery
 - about 25
 - and big3d agent 25
 - and gtmd agent 25
 - and probes 33
 - viewing statistics about connections 26
 - viewing status of connections 26

K

- Kilobyte/Second load balancing method, about 17

L

- last resort pool
 - about 38
 - assigning to a wide IP 38
- LDNS, and communication with the big3d agent 33
- LDNS entries, and state of BIG-IP DNS 35
- Least Connections load balancing method, about 17
- limit settings, defined 14
- links
 - and monitoring of outbound traffic 42
 - defined 41
- listeners
 - about wildcard 32
 - and network traffic 31
 - defined 31
- load balancing
 - about pool-level 16
 - about Topology 16
 - about wide IPs and pool order 14
 - and limit settings 14
 - and resource availability 13
 - and virtual server dependency 13
 - using tiered 9
 - using wide IP-level 9

- load balancing methods
 - and dynamic ratio [22](#)
 - Completion Rate [18](#)
 - CPU [18](#)
 - customizing QoS [21](#)
 - Drop Packet [17](#)
 - dynamic [11](#)
 - Fallback IP [19](#)
 - Global Availability [14](#)
 - Hops [18](#)
 - Kilobyte/Second [17](#)
 - Least Connections [17](#)
 - None [19](#)
 - Packet Rate [17](#)
 - QoS [19](#)
 - Ratio [15](#)
 - Return to DNS [18](#)
 - Round Robin [16](#)
 - Round Trip Times [17](#)
 - static [9](#)
 - Static Persist [18](#)
 - understanding QoS equation [19](#)
 - Virtual Server Capacity [17](#)
 - Virtual Server Score [17](#)
- load balancing process
 - about Prober pool status [33](#)
 - and Prober pools [32](#)
- logical network components
 - about distributed applications [42](#)
 - about pools [41](#)
 - about wide IPs [36](#)

M

- manual resume feature
 - and pools [22](#)
 - and virtual servers [23](#)
- monitor timeout, and virtual server status [41](#)

N

- named.conf
 - configuring using ZoneRunner [43](#)
 - defined [43](#)
- network traffic
 - listeners [31](#)
- None load balancing method, about [19](#)
- NTP servers, and BIG-IP DNS synchronization groups [29](#)

P

- Packet Rate load balancing method, about [17](#)
- persistent connections
 - about [37](#)
 - configuring [37](#)
- physical network components
 - about virtual servers [40](#)
 - and links [41](#)
 - defining servers [39](#)
- pool-level load balancing, about [16](#)
- pool members, and pools [41](#)

- pools
 - and CNAME records [41](#)
 - and pool member weighting [15](#)
 - and the manual resume feature [22](#)
 - defined [41](#)
 - organizing within wide IPs [14](#)
 - restoring availability manually [22](#)
 - weighting in wide IPs [15](#)
- preferred load balancing method, using when pool metrics are unavailable [22](#)
- probe responsibility, and BIG-IP systems [35](#)
- Prober pools
 - about [32](#)
 - about statistics [33](#)
 - about status [33](#)
 - and upgrading to version 11.x [34](#)
- probes, about [33](#)
- protocols, supported by BIG-IP DNS [50](#)

Q

- QoS equation
 - [21](#)
 - about customizing [20](#)
 - understanding [19](#)
- QoS method
 - about [19](#)
 - customizing equation [21](#)

R

- Ratio load balancing method
 - about [15](#)
 - and pool member weighting [15](#)
 - and pool weighting in wide IPs [15](#)
- resource availability, and load balancing [13](#)
- Return to DNS load balancing method, about [18](#)
- Round Robin load balancing method, about [16](#)
- Round Trip Times load balancing method, about [17](#)

S

- servers
 - about [39](#)
 - about third-party hosts [39](#)
 - about third-party load balancing [40](#)
 - as pool members [42](#)
- SSL, and iQuery [25](#)
- static load balancing methods [9](#)
- Static Persist load balancing method
 - about [18](#)
- statistics, and Prober pools [33](#)
- status, and Prober pools [33](#)
- synchronization
 - about [29](#)
 - and NTP servers [29](#)
- system upgrades, and Prober pools [34](#)

T

- tiered load balancing [9](#)

Topology load balancing
about 16

U

upgrades, and Prober pools 34

V

verifying virtual server availability, and fallback load balancing method 19

views

creating for DNS in ZoneRunner 46

defined 46

virtual server availability, verifying 19

Virtual Server Capacity load balancing method, about 17

virtual server dependency, and load balancing 13

virtual servers

and configuring dependencies 13, 40

and weighting of pool members 15

as pool members 41

configuring status dependency 41

defined 40

restoring availability manually 23

Virtual Server Score load balancing method, about 17

virtual server status, setting for clusters 41

W

wide IP-level load balancing 9

wide IPs

and distributed applications 42

and last resort pool 38

and persistent connections 37

and pool order 14

and pool weighting 15

and wildcard characters in names 36

assigning a last resort pool 38

defined 36

wildcard characters, and wide IP names 36

wildcard listeners, defined 32

Z

ZoneRunner

about 43

and configuring a hint zone 44

and configuring a zone 43

and configuring named 43

and creating DNS views 46

and viewing DNSSEC records 36, 50

zones

configuring hint 44

configuring using ZoneRunner 43

zone transfers, and BIG-IP DNS 44

