Edgecast

Route (DNS) Administration Guide

edgecast

Disclaimer

Care was taken in the creation of this guide. However, Edgecast cannot accept any responsibility for errors or omissions. There are no warranties, expressed or implied, including the warranty of merchantability or fitness for a particular purpose, accompanying this product.

Trademark Information

About This Guide

Route (DNS) Administration Guide
Version 2.41
11/22/2021

Table of Contents

Route

Introduction

The Route solution provides a variety of offerings through which you can achieve reliable, high performance, and secure DNS service. This solution allows you to perform all of the following:

- Manage a DNS zone.
- Set up a secondary DNS zone.
- Load balance traffic across servers.
- Establish a failover system that prevents an interruption to site traffic when server(s) can no longer properly fulfill requests.
- Verify a server's capability to fulfill requests through health checks performed from around the world.
- Mitigate Distributed Denial of Service (DDoS) attacks on your servers by leveraging our distributed DNS network and technologies.

Scope

The purpose of this document is to provide information on how you can define your DNS configuration to accomplish the above. A brief description of each chapter can be found below.

Chapter	Description
DNS Zone Management	Describes how to administer a DNS zone and its records.
Server Health Checks	Describes what global health checks are and how to configure them to validate a server's capability to serve traffic.
DNS Load Balancing	Describes how a load balancing configuration can distribute requests among various servers. It also includes information on how a server outage affects the proportion of traffic served to the remaining healthy servers.
DNS Failover	Describes how to define a failover system that will divert traffic to a backup server when the primary server is no longer capable of serving traffic.
Secondary Zone	Describes how to set up a secondary zone. This type of setup allows our name servers to provide an authoritative answer for all zones transferred to our DNS service.

Module Comparison

The Route solution consists of the following modules:

- Managed (Primary) or Secondary DNS
- DNS Health Checks

Managed (Primary) or Secondary DNS Module

This module allows:

- DNS zone management capabilities through which various types of records (e.g., A, AAAA, CNAME, etc.) may be added to a zone.
- Secondary zone capabilities that allow the transfer of a zone from a master name server to our name servers.
- Load balance or failover traffic across multiple servers.

This module contains the following components:

Component	Description
Standard Routing	Identifies zones that do not leverage our load balancing/failover capabilities or advanced policies.
Adaptive Availability	This component allows traffic management (i.e., load balancing or failover) for: ▪ Address or CNAME records of a primary zone hosted by our Route solution. ▪ A CNAME record or subdomain of a primary zone hosted by another DNS service provider. **Note:** This module excludes zones on which advanced policies have been applied. Once an advanced policy has been applied to a zone, it falls under the Advanced Policy Routing component.
Advanced Policy Routing	Identifies zones that leverage advanced policies that allow routing based on IP address, country, ASN, IP groups, network groups, POP, or IP type. This type of zone can also leverage our load balancing/failover capabilities. **Note:** Advanced policy setup requires the engagement of our Professional Services team.

DNS Health Checks Module

The health of a server/domain that is part of a load balancing or failover configuration can be probed by polling it with an HTTP or TCP request. If the majority of our servers deem it to be unfit to serve traffic, then this module will take the following action:

- **Load Balancing:** If a server/domain is deemed unhealthy, it will be pulled out of load balancing and traffic will be redistributed proportionally to the remaining healthy servers/domains.
- **Failover:** If the primary server/domain is deemed unhealthy, all traffic will be switched to the backup server/domain.

Billing Activation

Our Route solution has been designed for ease of use. This means that the various modules offered in this solution can be configured and managed from a single location. It also means that you can quickly and easily activate a module without having to request it from your DNS account manager. The activation process for each module is described below.

Module	Activation Process
Managed (Primary) or Secondary DNS	A zone is considered billable once the zone has been queried. **Note:** A zone will be billed at the Adaptive Availability rate once a load balancing or failover configuration has been created.
Health Checks	A health check configuration is considered billable once it has been created.

Once a module has been activated, a recurring monthly charge may be assessed to your account.

Note: For more detailed billing information, please contact your DNS account manager.

Note: A load balancing, failover, and health check configuration can be performed from within a zone. This capability is provided as a matter of convenience. Adding a load balancing, failover, or health check configuration from within a zone will result in additional charges on your account.

Terminology

Definitions for common terms used in this document are provided below.

Term	Definition
Authoritative Name Server	Identifies a name server that contains all records for a specific DNS zone and is trusted to be accurate. A recursive name server can forward DNS requests to this type of name server.
Domain	The purpose of a domain is to identify a resource (e.g., computer) on the Internet. A domain is organized into a hierarchy. The following illustration demonstrates this hierarchy. 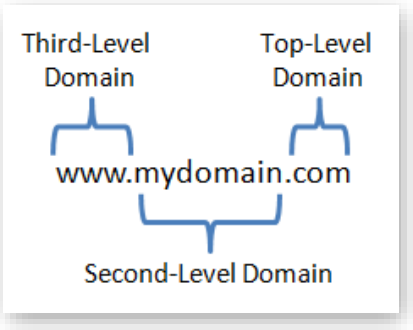 The above hierarchy is critical when resolving a domain, since the top-level domain and subsequent domains will be resolved by different authoritative name servers.
Domain Name System (DNS)	A hierarchical and distributed naming system for any resource connected to the Internet. This system includes the capability to translate domains into IP addresses.
Failover	Indicates the capability to redirect all traffic to a backup server/domain when the primary one is deemed unfit to resolve requests.
Health Check	The act of monitoring health status by requiring that a server acknowledge either a HTTP/HTTPS GET/POST, TCP, or TCP SSL request.
Hostname	A hostname is a label (e.g., www) used to identify a device.

Term	Definition
IP Address	Identifies a device (e.g., computer) by a numerical value. This document deals with two versions of IP addresses, which are: • **IPv4:** Most Internet traffic is routed through IPv4. ▪ **Example:** 101.10.10.253 • **IPv6:** A new version of IP address designed to address IPv4 address exhaustion. ▪ **Example:** 2001:0db8:85a3:0042:1000:8a2e:0370:7334
Load Balance	Indicates the capability to distribute requests between multiple servers or domains.
Master Name Server	Identifies a name server that is authoritative for the zone in question.
Name Server	A name server, or a domain name server, is a server that can provide an answer to DNS queries according to the information contained within a zone. A single name server can be authoritative for multiple zones. Two types of name servers are authoritative and recursive.
Record	A DNS zone can contain multiple records. A record defines the set of information through which a name server can provide an answer to a DNS request. We support the following records: A, AAAA, CNAME, MX, NS, SOA, SPF, SRV, and TXT.
Recursive Name Server	Identifies a name server that has the capability to cache the answer to a DNS query provided by an authoritative name server.
Root Name Server	Identifies a name server that can return the authoritative name server for each top-level domain.
Subdomain	Indicates the relative relationship between two domains. This relationship is explored below for a sample domain called "us.mydomain.com." • "us.mydomain.com" is a subdomain of the "mydomain.com" domain. • Both "us.mydomain.com" and "mydomain.com" are considered to be domains as well. **Note:** Each domain, with the exception of the top-level domain, is a subdomain as well as a domain.

Term	Definition
TSIG	Transaction Signature. Our solution allows the use of TSIG to provide a cryptographically secure method through which zone transfers may take place between a master and slave name server.
Zone	A zone consists of the portion of the DNS namespace over which authority has been delegated. It contains information (i.e., records) through which an authoritative name server can provide an answer to DNS queries.

Primary Zone Management

Overview

The Managed (Primary) or Secondary DNS module allows the creation and management of zones. A zone defines a set of data through which our authoritative name servers can provide a response to DNS queries. This data can be found in the records associated with your zone. In addition to record administration, zone management also allows the definition of load balancing and failover configurations for address and CNAME records within that zone.

Reminder: This section covers zone management and its relationship with load balancing and failover. Adding traffic management capabilities to your zone will incur additional charges on your account. Please contact your DNS account manager for more information.

Zones & Domains

Before setting up a zone, it is important to make a distinction between a zone and a domain. A domain is a generic label used to identify a resource (e.g., computer) on the Internet. A zone is a delegation of authority over a particular namespace. This delegation allows our name servers to be authoritative for that zone. This means that recursive name servers can only receive an authoritative answer for the domain associated with your zone through our authoritative name servers.

Important: A subdomain can be delegated to a different DNS service provider. For example, authority over "example.com" and "cdn.example.com" can be delegated to two different DNS service providers. However, this scenario requires that requests to "cdn.example.com" first be resolved by the name server that is authoritative for "example.com."

Note: A recursive name server can cache the response provided by our authoritative name servers. The length of time for this cached response is determined by the time-to-live (TTL) defined for the corresponding record in your zone.

The following scenario provides a high-level overview on how a web browser's query is resolved. It assumes that a zone has been previously created and configured on our DNS service.

1. A client uses a web browser to submit a request for "mydomain.com."
2. The web browser forwards this DNS request to a recursive DNS server as defined by the client's network configuration.
3. The recursive DNS server will check whether it is authoritative for the corresponding zone or if it previously has cached a response. If either of those conditions is true, it will provide an immediate response to the client.
4. Otherwise, the recursive DNS server will forward the request to a root name server. A root name server has a list of name servers that are authoritative for each top-level domain (e.g., COM, ORG, INFO, etc.).
5. The root name server will return the IP address of a name server that is authoritative for the requested domain's top-level domain (e.g., COM).
6. The recursive DNS server will then forward the DNS query to top-level name servers. These name servers have a list of name servers that are authoritative for each second-level domain (e.g., mydomain.com) associated with that top-level domain.
7. A top-level name server will then return the IP address of our authoritative name server.
8. The recursive DNS server will then forward the DNS query to our authoritative name server.
9. Our authoritative name server will provide an answer to the DNS query according to the records associated with your zone. For the purposes of this example, it will return the IP address defined in an A record.
10. The recursive DNS server will then relay the authoritative name server's answer to the client's computer and oftentimes cache it for future requests for the length of time associated with the TTL.
11. The client's user agent will then handle the request according to that answer.

The above scenario is illustrated below.

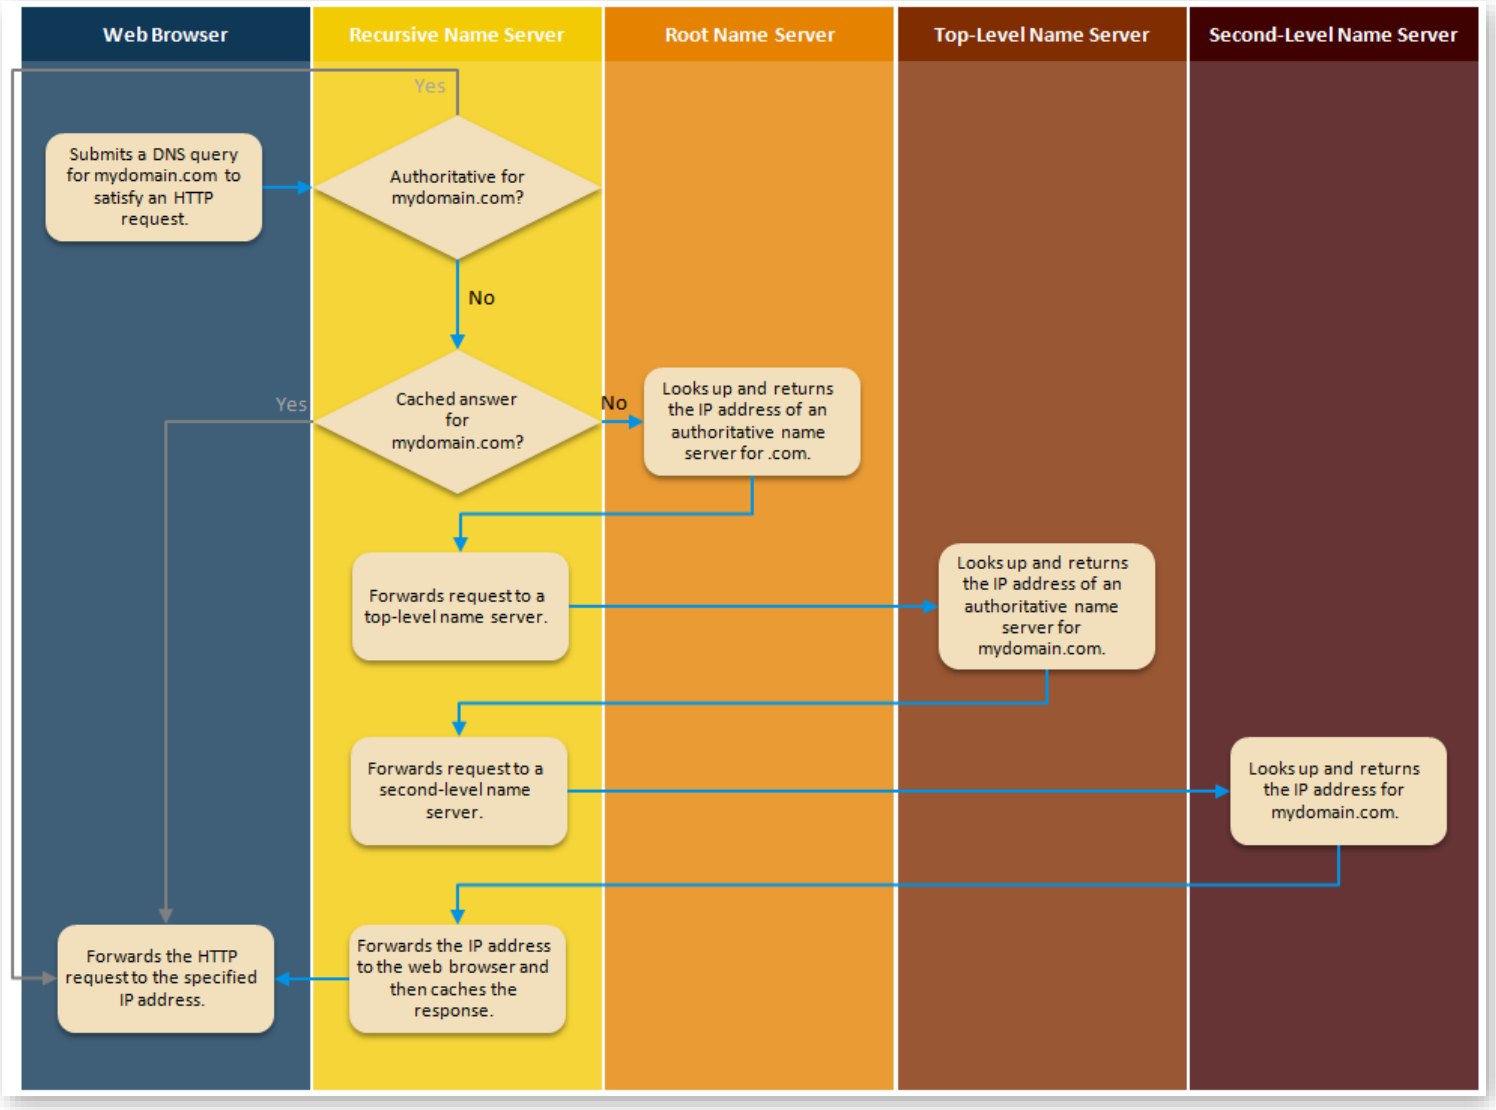

Resolving a Web Browser's DNS Query for mydomain.com

As previously mentioned, a name server is only authoritative for the zones associated with it. In the above example, the "mydomain.com" zone was delegated to our DNS service. In the following illustration, we will see how DNS is resolved when "us.mydomain.com" is delegated to our DNS service and the "mydomain.com" zone remains with a third-party DNS service provider.

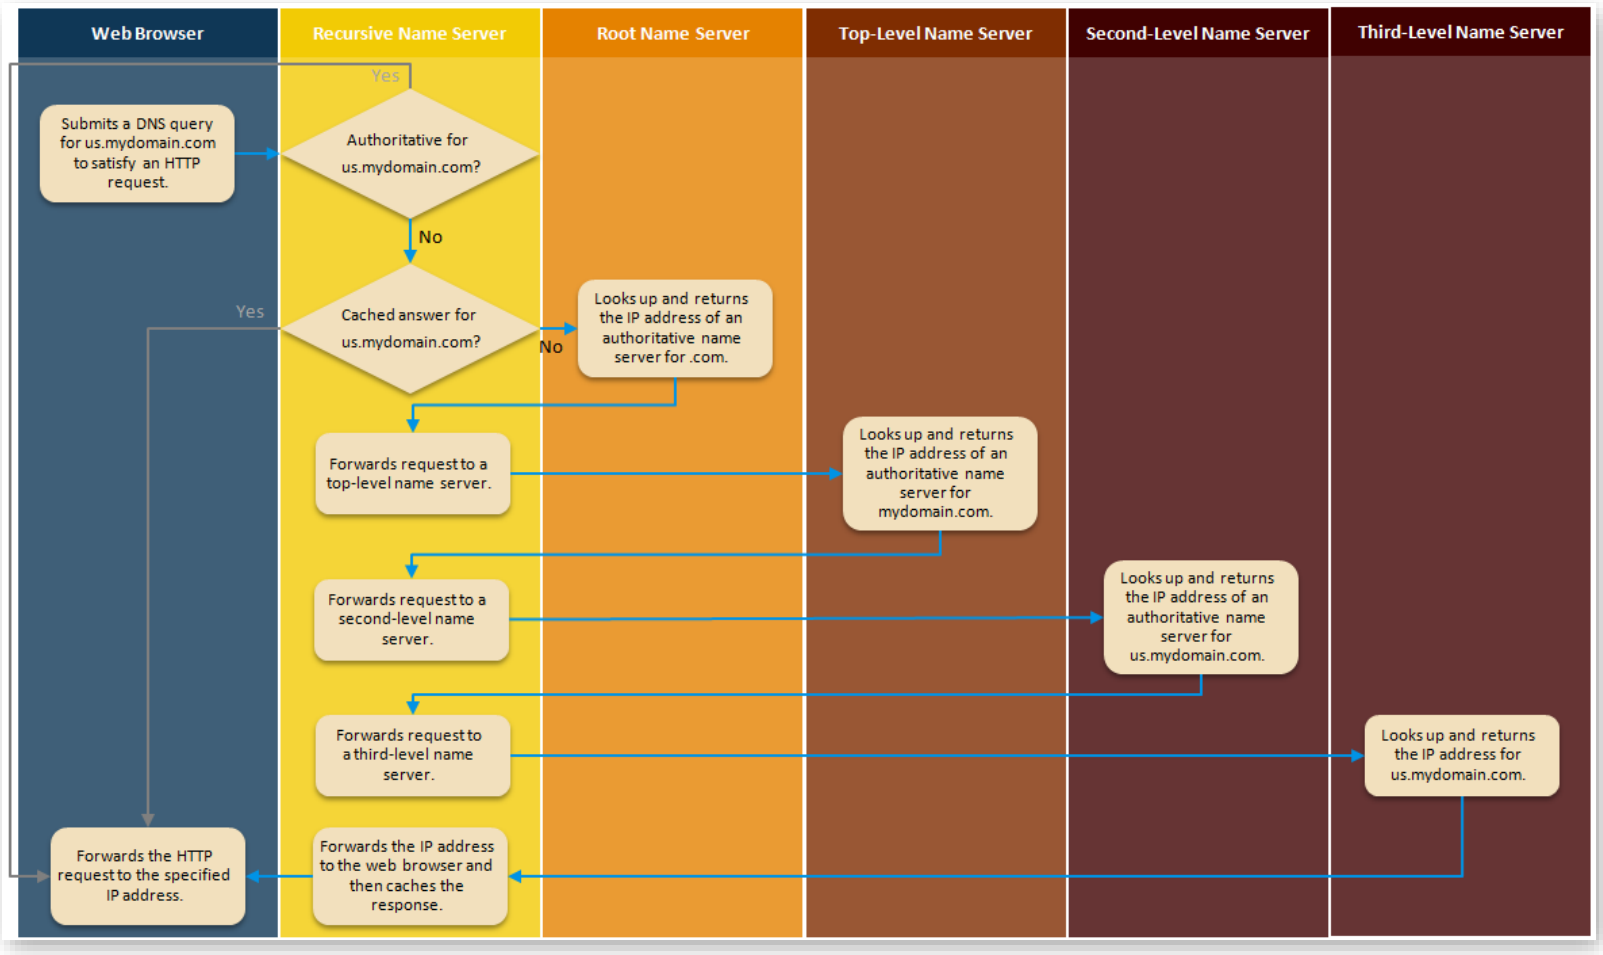

Resolving a Web Browser's DNS Query for us.mydomain.com

It is important to note the following:

- An authoritative answer for a domain can only be provided by a name server that has authority over the corresponding zone. In the above example, "us.mydomain.com" has been delegated to our service. Although it is a subdomain of "mydomain.com," an authoritative answer can only be provided by our name servers. Likewise, an authoritative answer for "mydomain.com" can only be provided by the corresponding second-level name servers.
- Although authority over "us.mydomain.com" has been delegated to our service in this scenario, we are still reliant on a third-party second-level name server to provide our IP address to the recursive DNS server. If the third-party DNS service that owns those second-level name servers is unavailable or non-responsive, then the DNS query will be unable to reach our authoritative name servers. Therefore, we highly recommend that you delegate your second-level domain to our service by creating a zone for it and then point your domain to our name servers. This will lead to a faster response and avoid potential issues that may arise when a third-party DNS service is used.

Zone Components

A zone consists of the following components:

- **Zone name:** Assigns a label to a zone. This label also defines the origin (e.g., mydomain.com) associated with a zone.
- **Comment:** Associates a comment or a description with a zone. This option is solely provided for informational purposes and does not affect your zone configuration.
- **Status:** Indicates whether a zone is active or inactive. Our authoritative name servers will only leverage active zones when resolving DNS queries.
- **Records:** A set of records should be associated with a zone. Each record provides information through which our name servers can provide answers to DNS queries. A description for each record type is provided in the **Record Administration** section below.
- **Load Balancing:** Defines a load balancing group that determines how traffic is distributed to your servers. Keep in mind that the configuration of this feature will result in an additional charge.
- **Failover:** Defines a failover group that consists of a primary and backup server/domain. All traffic will be served to the primary server/domain until it is unhealthy. At which point, all traffic will be redirected to the backup one. Keep in mind that the configuration of this feature will result in an additional charge.

DNS Zone Administration

This section describes how to administer zones and its components.

Zone Creation/Modification

Keep the following information in mind when creating or modifying a zone:

- Upon creating a zone, SOA, NS, A, and AAAA records will be automatically added to it. These records are required to use our DNS service and cannot be modified. The NS and host records are used to define vanity name servers for your zone. Vanity name servers lend a professional appearance to your site's DNS.
- In addition to the above mandatory records, you should create additional records that define how DNS queries will be handled. For information on how to manage records, please refer to the **Record Administration** section below.
- The maximum length of a zone name is 230 characters.
- Once a zone has been created, its name (e.g., mydomain.com) cannot be modified.
- A zone may be created by importing a zone file.
    - Only zone files exported using the BIND format may be imported into our system.
    - Only import zone files that contain a single zone.
    - Once the zone has been imported, verify that all of the desired records have been transferred.
- The **Status** option determines whether our name servers will provide an answer for the corresponding zone.
- The creation or modification of a zone will not be saved until **Submit Group** is clicked.
- Once a zone is ready for Production traffic, delegate the corresponding domain through a domain registrar. For more information, please refer to the **Switching DNS Service Provider** section below.

To create a zone

1. Navigate to the **Route (DNS)** page.
2. Click **Add New**.
3. In the **Type** option, select "Zone."
4. In the **Name** option, type the name of the zone (e.g., mydomain.com) for which our name servers will be authoritative.
5. Optional. Import a zone that was previously exported from a third-party DNS service provider. Click **Browse**, browse to and select the desired zone file, and then click **Open**.
6. Use the **Status** option to indicate whether our name servers will provide an answer for this zone.
7. Click **Add**.
8. Add the desired records. For more information, please refer to the **Record Administration** section below.
9. Click **Submit Group** to save your changes.

To modify a zone

1. Navigate to the **Route (DNS)** page.
2. Click on the desired zone.
3. Optional. Set a comment by clicking **Edit** which appears to the right of the desired zone. Define the desired comment and then click **Done**.
4. Optional. Change the zone's status by clicking **Edit** which appears to the right of the desired zone. Select the desired status and then click **Done**.
5. Add the desired records. For more information, please refer to the **Record Administration** section below.
6. Click **Submit Group** to save your changes.

Zone Deletion

A zone can be deleted from the **Route (DNS)** page by simply hovering over the desired zone and then clicking  . When prompted, confirm the deletion of the zone.

Key information:

- Zone deletions are permanent. Once a zone has been deleted, it cannot be recovered.
- An alternative to zone deletion is to temporarily disable it by changing its status to "Inactive." Learn more.

Record Administration

Records determine how DNS queries will be resolved by our authoritative name servers. The types of records that can be associated with a zone are described below.

Name	Type	Common Usage	Configuration
A	Address	Maps a hostname to an IPv4 address.	**Name** Identifies a hostname (e.g., @ or www). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the IPv4 address that will be mapped to the specified domain.
AAAA	Address	Maps a hostname to an IPv6 address.	**Name** Identifies a hostname (e.g., @ or www). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the IPv6 address that will be mapped to the specified domain.
CAA	Certificate Authority Authorization	Defines the Certificate Authorities (CAs) that are authorized to issue certificates for the domain corresponding to this zone.	**Name** Identifies a hostname (e.g., @ or www). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Defines your CAA policy for the specified hostname using the following syntax: *flags* issue|issuewidl|iodef *value* Learn more (Wikipedia). Learn more (RFC 6844). **Sample value:** 0 issue ca.example.net

Name	Type	Common Usage	Configuration
CNAME	Alias	A Canonical Name record maps a hostname to another hostname or FQDN.	**Name** Identifies a hostname (e.g., @ or www). **Note:** Set the **Name** option to the "@" symbol to point the CNAME record to the zone apex (e.g., example.com). **Note:** An asterisk may be used as a "starts with" wildcard. However, it may not be used in the middle or at the end of the specified value. **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the domain that will be mapped to the hostname identified by the **Name** option. A period must be appended to this domain.
MX	Mail Exchange	Maps a hostname to a mail server. Indicates the SMTP gateways to which mail can be delivered.	**Name** This read-only option uses the @ symbol to identify the hostname associated with the zone. **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the mail server's priority and hostname. The format for this field is: • *Priority Hostname*. • **Example:** 0 smtp1.mydomain.com. Each of the above terms is defined below. • **Priority:** Represents an integer value that defines the mail server's priority. A mail server with a lower priority value is given preference over other mail servers defined for the same zone. • **Hostname:** Represents the mail server's hostname. **Note:** The specified hostname must end with a period (e.g., example.com.).

Name	Type	Common Usage	Configuration
NS	Name Server	Delegates a hostname to a name server.	**Name** Identifies a hostname (e.g., ns1). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the name server to which the specified hostname will be delegated.
PTR	Pointer	Maps an IPv4 address to a hostname. Use this type of record when setting up a reverse DNS lookup.	**Name** Identifies an IPv4 address by its fourth octet (e.g., 100). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies a hostname that will be mapped to the specified IPv4 address. **Note:** The specified hostname must end with a period (e.g., example.com.).
SOA	Start of Authority	Provides authoritative information about a DNS zone.	This record is automatically defined when you create a zone and it cannot be modified.
SPF	Sender Policy Framework	Defines the mail servers that can legitimately send e-mails from the zone's domain.	**Name** Identifies the name of the zone. This option is read-only. **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the mail server's address.
SRV	Service Locator	Identifies the location of a service (e.g., FTP).	**Name** Identifies the name and protocol associated with the service. You may also append a hostname to this value (e.g., _http._tcp.mydomain.com). Make sure that the name starts with an underscore (e.g., _tcp, _http, or _udp) and use a period to separate multiple protocols (e.g.,

Name	Type	Common Usage	Configuration
			_tcp._http). **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** Identifies the service's priority, weight, port, and target. These properties should be specified as indicated below. • *Priority Weight Port Target* Each of the above terms is defined below. • **Priority:** An integer that defines the service's priority. A service with a lower priority value is given preference over other services with the same protocol and target. • **Weight:** An integer that determines a service's contact order when there are other services with the same protocol, target, and priority. • **Port:** Represents the service's port number (e.g., 80). • **Target:** Represents the service's hostname (e.g., ftp.mydomain.com).
TXT	Text	This record allows text to be associated with a zone. Among its many uses, it can store SPF data.	**Name** This field can be set to free form text. **TTL** Indicates the length of time (in seconds) that a server should cache the record. **Value** This field can be set to free form text.

Defining a Record's Hostname

Certain records (e.g., A or AAAA records) require a hostname (e.g., www or us) for the **Name** option. Keep in mind that the origin associated with a zone (e.g., mydomain.com) will be appended to the specified host label. For example, if you specify "us" as the host label in the **Name** option, then the hostname corresponding to that record would be "us.mydomain.com." Additional syntax information is provided below.

Points To	Syntax	Description
Origin	@	The @ symbol identifies the origin (e.g., mydomain.com) associated with the current zone. **Note:** The proper usage of this syntax is to set the **Name** option to the @ symbol. No other characters should be specified.
Host Label	*Host Label*	Specify the desired host label in the **Name** option. For the purposes of our DNS service, it will treat the record as if the origin (e.g., mydomain.com) had been appended to it. **Tip:** Do not specify a hostname (e.g., www.mydomain.com).

System-Defined Mandatory Records

Upon creating a zone, a set of records required by our DNS service are automatically added to it. These system-defined records cannot be modified or deleted. A list is provided below.

- **SOA:** Defines the primary name server associated with your zone.
- **NS:** Defines the name servers to which queries will be directed. These name servers are known as vanity name servers.
- **A:** Defines the IP addresses (IPv4) corresponding to our vanity name servers.
- **AAAA:** Defines the IP addresses (IPv6) corresponding to our vanity name servers.

Note: Although the above records are mandatory and cannot be modified or deleted, you may always add, modify, and delete, as needed, additional NS, A, and AAAA records to your zone.

Zone Apex Support for CNAME Records

A CNAME record can be set to a zone apex (aka naked domain or root domain). This allows the zone apex (e.g., example.com) to resolve to a subdomain (e.g., www.example.com).

Configuration

Point a CNAME record to the zone apex (e.g., example.com) through the following configuration:

Setting	Value	Description
Name	@	The @ symbol indicates that the CNAME record should point to the zone apex.
TTL	*{Seconds}*	Set the TTL to the length of time (in seconds) that a DNS server should cache the record.
Value	*{Hostname}*	The specified hostname should point to an A or AAAA record.

How Does It Work?

Our authoritative name servers resolve a CNAME record to an A or AAAA record. As a result, the hostname defined in the CNAME record will not be delivered to the requester. Rather, our authoritative name servers will serve the IP address associated with the A or AAAA record.

Sample Scenario

This scenario explores how a CNAME record that points to the zone apex is handled by our DNS service. In this example, a zone called example.com contains the following CNAME and A records:

CNAME record configuration:

Setting	Value
Name	@
TTL	3600
Value	www.example.com.

A record configuration:

Setting	Value
Name	www
TTL	3600
Value	192.0.2.100

Requests that point to example.com will be resolved by our authoritative DNS servers to 192.0.2.100. The requester will be unaware that example.com actually points to www.example.com.

Reverse DNS Lookup (PTR Records)

A reverse DNS lookup identifies the hostname associated with an IPv4 address by leveraging a pointer (PTR) record.

Configuration

Set up a reverse DNS lookup by performing the following steps:

1. Create a zone within the in-addr.arpa domain that is specific to the first three octets of the desired IPv4 address. Reverse the order of the first three octets when setting up your zone.

   Syntax:

   {3rd Octet}.{2nd Octet}.{1st Octet}.in-addr.arpa

   Example:

   Create the following zone for 192.0.2.100:

   2.0.192.in-addr.arpa

2. Create a PTR record for the desired IP address within the above zone.

Setting	Value	Description
Name	*{Fourth Octet}*	Set to the fourth octet of the desired IPv4 address.
TTL	*{Seconds}*	Set the TTL to the length of time (in seconds) that a DNS server should cache the record.
Value	*{Hostname}*	Indicate the hostname associated with this IPv4 address.

   Example:

   Create a PTR record for 192.0.2.100 by setting the name to "100," TTL to "3600," and the value to the desired hostname (e.g., www.example.com.).

Upon setting up the above configuration, a reverse DNS lookup for 192.0.2.100 returns:

100.2.0.192.in-addr.arpa. 3600 IN PTR www.example.com.

Record Creation

A record can be created by simply performing the following steps:

1. Navigate to the **Route (DNS)** page.
2. Click on the desired zone.
3. Click **Add Record**.
4. In the **Type** option, select the desired type of record.
5. Specify the desired properties of that record.
6. Click **Add**.
7. Click **Submit Group** to save your changes.

Record Modification

A record can be modified by simply performing the following steps:

1. Navigate to the **Route (DNS)** page.
2. Click on the desired zone.
3. Click **Edit** next to the record that you would like to modify.
4. Make the desired modifications.
5. Click **Done**.
6. Click **Submit Group** to save your changes.

Reminder: The default set of system-defined records cannot be modified.

Record Deletion

A record can be deleted by simply performing the following steps:

1. Navigate to the **Route (DNS)** page.
2. Click on the desired zone.
3. Click **Edit** next to the record that you would like to delete.
4. Click  .
5. Click **Submit Group** to save your changes.

Note: The  icon will only appear next to a record once it has been saved. If a new record has yet to be saved, then you can either discard all of your changes by navigating to a different page or save your zone and then delete the record using the above instructions.

Reminder: The default set of system-defined records cannot be deleted.

Zone Status

A zone's status determines whether DNS service will be enabled for that zone. Common usages for this feature are:

- Preventing our name servers from providing answers for a zone that has not been properly configured. Setting a new zone's status to inactive allows you to create the desired records and double-check your configuration before enabling our DNS service.
- Testing your configuration. Once you are satisfied that your zone has been configured properly, you can disable it until you want to start serving production traffic on it.

Keep the following information in mind with regards to status and billing:

- A zone's status does not affect whether it is billable. Once traffic has flowed through a zone, it has been activated and will be considered a billable zone even if it never receives additional traffic.
- Only billable zones are included when calculating the total number of zones associated with your account.
- Only by deleting a billable zone can it be excluded from your total zone calculation. For detailed information on billing, please contact your DNS account manager.

Switching DNS Service Provider

Creating a zone will not automatically redirect DNS traffic to our Route name servers. Once a zone has been properly configured for production traffic, its traffic needs to be directed to our Route name servers by changing the name servers at your domain registrar. This may require the creation of glue records, which are also known as Host or Name Server records. This type of record should identify either of the following:

- Vanity name servers
- Route-branded name servers

Vanity Name Servers

Upon creating a primary or secondary zone, our Route solution automatically creates NS records for each of our vanity name servers. Delegate your primary or secondary zone by pointing your glue records to our vanity name servers. The naming convention for our vanity name servers varies according to whether you are delegating a primary or secondary zone.

Zone Type	Naming Convention	Example
Primary	ns#.<Zone Name>.	ns1.example.com.
Secondary	s#ns#.<Zone Name>.	s1ns1.example.com.

In the above naming convention, the pound symbol (#) represents a sequential number and <Zone Name> identifies the name of your zone.

We provide 4 vanity name servers and each one may fulfill IPv4 and IPv6 requests. As a result, our vanity name servers leverage 8 IP addresses to serve your zone. The following sample scenario provides a listing of the IP addresses in use by our vanity name servers.

Primary Zone Delegation Example

This example indicates the records that will be created for a primary zone called "example.com."

Record Type	Name	Value
NS	@	ns1.example.com.
NS	@	ns2.example.com.
NS	@	ns3.example.com.
NS	@	ns4.example.com.
A	ns1	192.16.16.5
A	ns2	192.16.16.6
A	ns3	198.7.29.5
A	ns4	198.7.29.6
AAAA	ns1	2606:2800:3::5
AAAA	ns2	2606:2800:3::6
AAAA	ns3	2606:2800:c::5
AAAA	ns4	2606:2800:c::6

Delegate this zone to our DNS service by adding all of our vanity name servers (e.g., ns1.example.com.) as custom name servers to your domain registrar.

Note: Domain delegation varies by domain registrar. Additional information can be found in the following Route Help Center article: FAQ: Route.

Route-Branded Name Servers

A primary zone may be delegated to our Route-branded name servers instead of vanity name servers.

Important: DNS testing tools may warn that stealth name servers have been detected. This warning is due to the presence of our vanity name servers within your zone. These vanity name servers cannot be modified or deleted.

To delegate a primary zone using Route-branded name servers

1. Define the following records within your zone:

Record Type	Name	Value
NS	@	ns1.edgecastdns.net.
NS	@	ns2.edgecastdns.net.
NS	@	ns3.edgecastdns.net.
NS	@	ns4.edgecastdns.net.
A	ns1	192.16.16.5
A	ns2	192.16.16.6
A	ns3	198.7.29.5
A	ns4	198.7.29.6
AAAA	ns1	2606:2800:3::5
AAAA	ns2	2606:2800:3::6
AAAA	ns3	2606:2800:c::5
AAAA	ns4	2606:2800:c::6

2. Delegate your domain by registering our Route-branded name servers (i.e., ns[1-4].edgecastdns.net.) with your domain registrar.

DNS Health Checks

Overview

A key to ensuring that your site's traffic flows properly is to check at regular intervals that your web servers can provide a response to requests. Our DNS Health Checks module is designed to check server/domain health status at regular intervals by sending an HTTP, HTTPS, TCP, or a TCP SSL request from our health checks agents. The worldwide distribution of our health check agents ensures that network latency doesn't result in a misdiagnosis of a server/domain's health state.

Reminder: Health checks can be applied to traffic management configurations (e.g., load balancing or failover). Adding traffic management capabilities to a zone will incur additional charges on your account. Please contact your DNS account manager for more information.

How Does It Work?

Our health check agents can be configured to poll a server at regular intervals using one of the following types of requests:

Request Type	Description
HTTP/HTTPS GET	Polls a server using a GET request over HTTP or HTTPS.
HTTP/HTTPS POST	Polls a server using a POST request over HTTP or HTTPS.
TCP/TCP SSL	Polls a server by opening a connection over TCP or TCP SSL.

Note: None of the above requests contain a request body.

A server/domain will be considered healthy if both of the following conditions are met:

- The response indicates that a valid connection was established (e.g., 200 OK).
- An additional check will be performed if content verification has been enabled on the health check. The body of the server's response to a health check request must contain a user-defined word or phrase.

Our health check agents are distributed around the world. Each of them will poll a monitored server/domain every few seconds, as determined by the health check configuration. A simple majority consensus will then be used to determine whether traffic should be pulled from a server/domain.

Note: A health check may fail due to a reason that is unrelated to health status (e.g., a connectivity issue on a single transit/peering route).

The following illustration demonstrates that a web server can be considered healthy even though a health check failed.

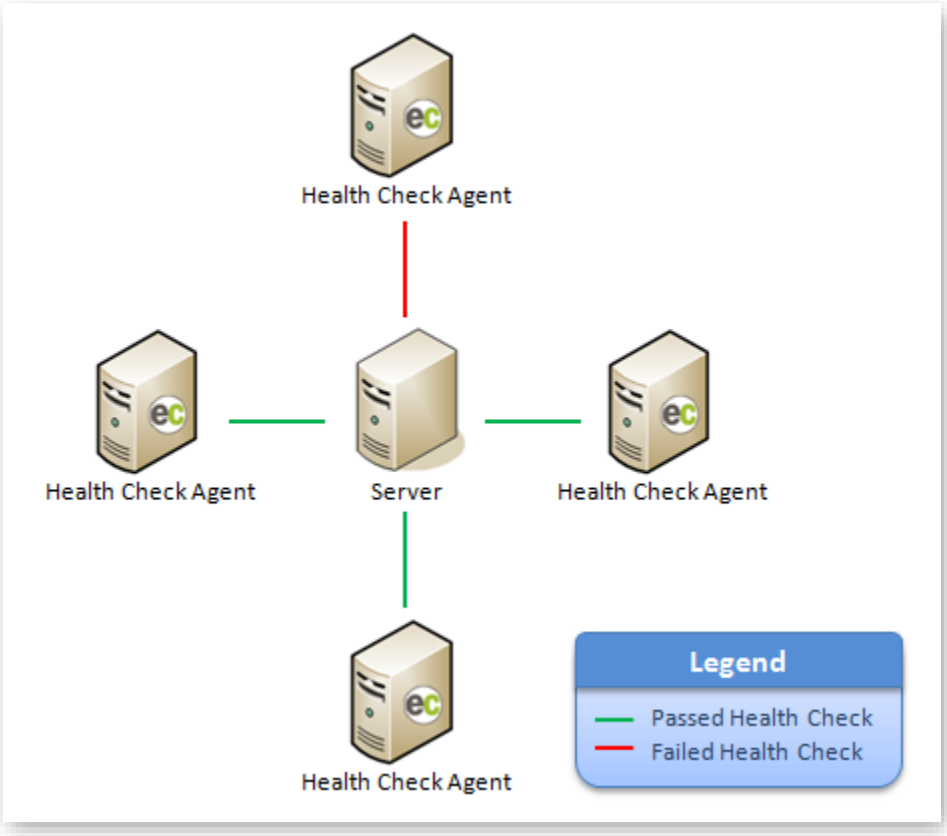

Health Check Consensus: Pass

A server/domain is considered healthy until a majority of health check agents indicate otherwise. In the following illustration, we see that there is a unanimous consensus that a server is unhealthy. This will cause our DNS service to stop serving traffic to that server as indicated below.

- **Load Balancing:** Traffic will be redistributed proportionally to the remaining servers in that load balancing group.
- **Failover:** If a primary server or domain is found to be unhealthy, then traffic will be redirected to the backup one.

Note: If all servers/domains in a load balancing or failover group have failed the majority of their health checks, then a DNS query to that group will be resolved to a default IP address associated with that group.

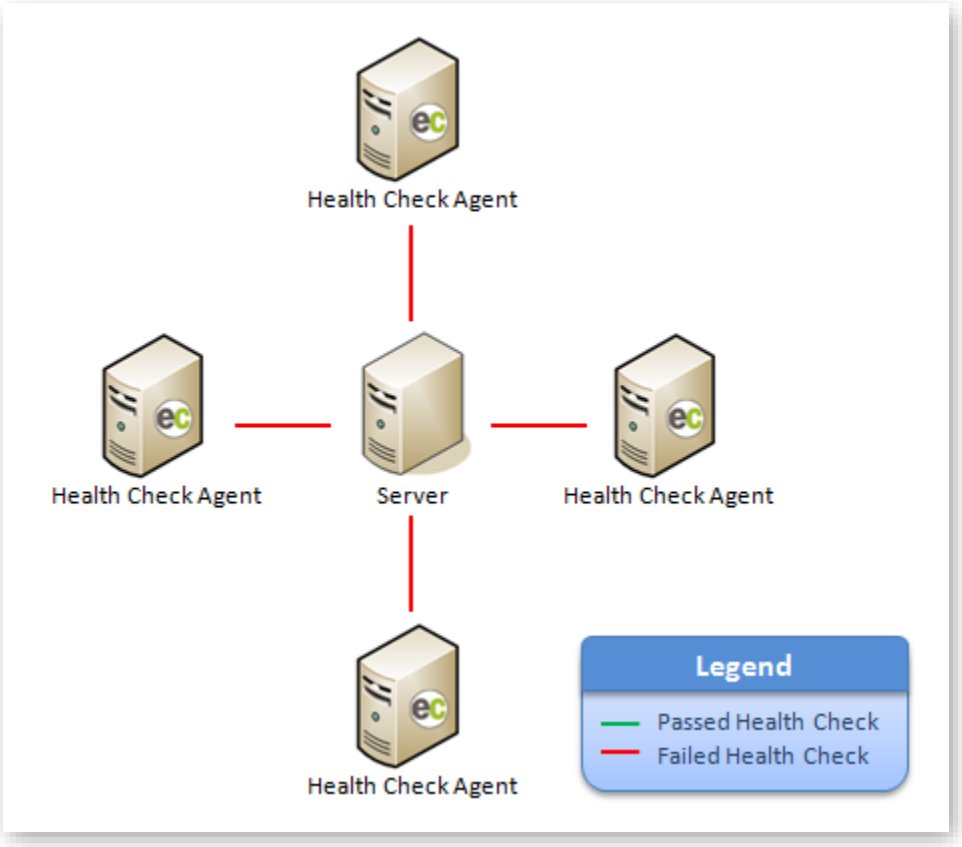

Health Check Consensus: Fail

Once traffic is pulled from a server/domain, your health check configuration determines whether it will be automatically or manually reintegrated. Automatic reintegration means that traffic will flow through that server/domain as soon as a consensus is established that it is healthy. Manual reintegration means that you have to manually update your load balancing or failover configuration to indicate that it is now ready to receive traffic.

Configuration

Keep the following information in mind when configuring a health check:

- Each health check configuration is specific to a single IP address/domain. As a result, the health check for each IP address or domain must be configured separately.
- An IP address/domain is eligible for health checks when it has been associated with a load balancing or failover configuration. For example, a health check configuration cannot be applied to an address record (i.e., A or AAAA) until it is part of a load balancing or failover configuration.
- The URI or IP address used to check a server/domain's health does not have point to it. This allows the use of a server that monitors web servers and reports health information. Content verification may be used to verify the health of an individual server/domain.
- The creation, modification, or deletion of a health check configuration will not be saved until **Submit Group** is clicked.

A sample health check configuration is shown below.

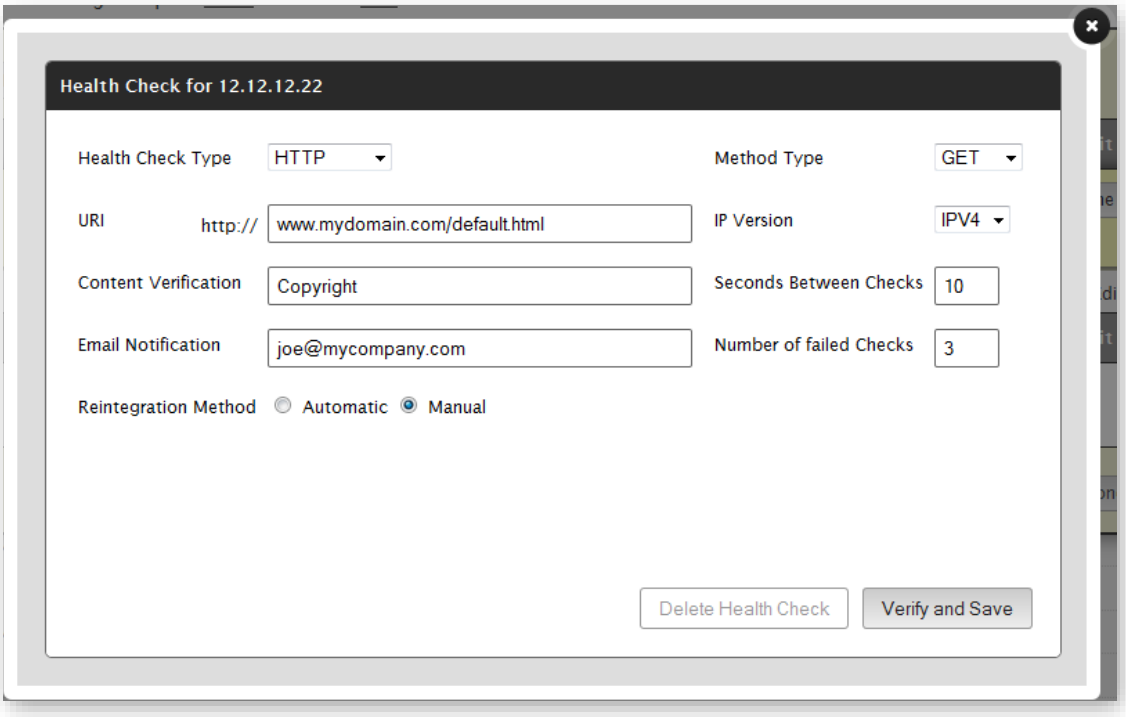

Health Check Configuration

The following table describes the various options that can be defined.

Option	Description
Health Check Type	Determines the type of health check that will be performed. The available options are: • **HTTP/HTTPS:** Use these health check types to query a web server using a GET or POST request over port 80 and 443, respectively. • **TCP/TCP SSL:** Use these health check types to establish a TCP connection to other services, such as FTP, SMTP, POP3, IMAP, etc.).
Method Type	**HTTP or HTTPS Health Checks Only** Determines whether a GET or a POST request will be used to poll health status. **Note:** Requests used to poll servers will not include a request body.
URI	**HTTP or HTTPS Health Checks Only** Determines the URI for the GET/POST request that our health check agents will use to poll health status. The URI's protocol (i.e., http:// or https://) was specified in the **Health Check Type** option and therefore should not be included in this option.
IP Version	**HTTP or HTTPS Health Checks Only** Determines whether the HTTP/HTTPS request should be performed using IPv4 or IPv6.
Ipv4 or IPv6	**TCP or TCP SSL Health Checks Only** Defines the IPv4 or IPv6 IP address to which the TCP request will be directed.
Port Number	**TCP or TCP SSL Health Checks Only** Determines the port over which the request will be performed. Make sure that monitored server(s) allow communication over this port.
Content Verification	Defines the set of ASCII characters through which a server's health will be verified. Our health check agents will look for the specified text in the body of the server's response to a health check request. If this option is left blank, then our servers will simply verify the following: • **HTTP:** The HTTP status code for the response is a 200 OK. • **TCP:** A connection could be established.
Seconds Between Checks	Determines the time interval, in seconds, between polling requests sent by our health check agents to monitored server(s).

Option	Description
Email Notification	Determines the e-mail address to which a notification will be sent whenever there is a change in the health state of the server/domain associated with this health check configuration.
Number of failed Checks	Each health check agent will attempt to communicate with a monitored server/domain to check whether it can pass a health check. A health check agent will record the results of each attempt (i.e., pass or fail), but it cannot make an assessment on health state until the same result occurs consecutively the number of times specified in this option. **Reminder:** Although each health check agent keeps track of whether a server/domain has passed/failed health checks, a majority consensus of all health check agents is required to change its health state.
Reintegration Method	Our DNS service will stop distributing traffic to a server/domain that has failed a majority of health checks. This option determines the process through which a server/domain can be reintegrated into a load balancing or failover group. Set this option to one of the following: • **Automatic:** Indicates that our health check agents will continue to poll a server/domain that has failed a majority of health checks. Once it is deemed healthy by a majority of our health check agents, it will automatically be reintegrated into your load balancing or failover configuration. The **Number of failed Checks** option determines the number of successful validations required before it can be considered healthy. • **Manual:** Indicates that manual intervention is required to reintegrate a server/domain into a load balancing or failover group. **Note:** Information on how to manually reintegrate a server can be found in the **DNS Load Balancing** and the **DNS Failover** chapters.

To add a health check configuration

1. Navigate to the **Route (DNS)** page.
2. Perform one of the following:
    - **Load Balancing/Failover (CNAME/Subdomain):** Click on the desired load balancing or failover configuration.
    - **Load Balancing/Failover (Zone):**
        i. Click on the desired zone.
        ii. Click **Edit** corresponding to the load balancing or failover configuration that you would like to modify.
3. Click **Edit** corresponding to the desired IP address or CNAME.
4. Click **Add HC**.
5. In the **Health Check Type** option, select the type of health check that will be used to poll a server/domain and then perform one of the following procedures:
    - **HTTP/HTTPS:**
        i. In the **Method Type** option, select whether a GET or POST request will be sent.
        ii. In the **URI** option, define the URI for the GET or POST request that will be sent.
        iii. In the **IP Version** option, select whether the request will be sent over IPv4 or IPv6.
    - **TCP:**
        i. In the **Ipv4 or IPv6** option, type the IP address to which the TCP request will be sent.
        ii. In the **Port Number** option, define the port over which TCP communication will take place.
6. Optional. In the **Content Verification** option, define the text through which health status will be verified. Our health check agents will search the body of the requested content for the specified text.
7. In the **Seconds Between Checks** option, define the time interval for which a health check will be performed from each of our health check agents.
8. In the **Number of failed Checks** option, define the number of consecutive times that a health check agent must achieve the same result (i.e., pass or fail) before it can determine health state (i.e., healthy or unhealthy).

9. In the **Email Notification** option, type the e-mail address to which notifications will be sent for unhealthy servers/domains.
10. In the **Reintegration Method** option, define whether the server/domain associated with this health check can automatically be reintegrated into your load balancing or failover configuration after recovering from failed health checks.
11. Click **Verify and Save**. The health check configuration will be verified by polling a server/domain according to the defined settings.
12. **Zone Only:** Click **Done**.
13. Click **Submit Group** to save your changes.

To modify a health check configuration

1. Navigate to the **Route (DNS)** page.
2. Perform one of the following:
    - **Load Balancing/Failover (CNAME/Subdomain):** Click on the desired load balancing or failover configuration.
    - **Load Balancing/Failover (Zone):**
        i. Click on the desired zone.
        ii. Click **Edit** corresponding to the desired load balancing or failover configuration.
3. Click **Edit** corresponding to the desired IP address or domain.
4. Click **Edit** under the **Health Check** column.
5. Make the desired changes.
6. Click **Verify and Save**. The health check configuration will be verified by polling a server/domain according to the defined settings.
7. **Zone Only:** Click **Done**.
8. Click **Submit Group** to save your changes.

To delete a health check configuration

1. Navigate to the **Route (DNS)** page.
2. Perform one of the following:
    - **Load Balancing/Failover (CNAME/Subdomain):** Click on the desired load balancing or failover group.
    - **Load Balancing/Failover (Zone):**
        i. Click on the desired zone.
        ii. Click **Edit** corresponding to the load balancing or failover configuration that you would like to modify.
3. Click **Edit** corresponding to the desired IP address.
4. Click **Edit** under the **Health Check** column.
5. Click **Delete Health Check**.
6. **Zone Only:** Click **Done**.
7. Click **Submit Group** to save your changes.

Health Check Status Information

An icon will appear next to each monitored server indicating its current status. A server's health check status is determined by a simple majority consensus. The following table provides a list of icons and their definitions.

Icon	Label	Description
	Success	Indicates that a majority of our health check agents have successfully verified that the monitored server/domain is healthy.
	Failed	Indicates that a server/domain failed a majority of its health checks. Our DNS service will not serve traffic to servers/domains in this state.
	Admin Down	Indicates that the server/domain has passed a majority of health checks, but traffic will not be directed to it as a result of a load balancing or failover configuration. The most common scenario for this state occurs when a server/domain configured for manual reintegration has recovered from a failed health check state but has not been reintegrated into a traffic management configuration.
	Unknown	Indicates that our health check agents are unaware of the server/domain's health state. This state occurs when a new health check configuration has been established. Once our health check agents make a decision on health status, a new status will be assigned to it.

Firewall Access and Monitored Servers

The DNS Health Checks module of our Route solution can monitor servers for a variety of DNS and performance-related factors. This type of monitoring requires each monitored server to allow access to our health check agents (i.e., servers). A list of the IP addresses (IPv4 and IPv6) corresponding to our agents is provided below.

IP blocks (IPv4):

```
5.104.64.0/24
46.22.76.128/25
108.161.247.0/24
152.195.0.0/16
152.199.103.0/24
152.199.111.0/24
192.16.2.0/24
192.16.57.0/24
192.16.60.0/24
198.7.18.0/24
198.7.22.0/24
```

IP blocks (IPv6):

```
2606:2800::/32
```

User Agent

The Health Check module polls servers and domains using the following user agent:

- EdgeCast-HealthCheckModule/1.0

Filter out requests that match this user agent when generating a site's traffic analytics.

DNS Load Balancing

Overview

The process of distributing traffic for a particular domain across multiple servers is known as load balancing. This distribution of requests ensures data availability through redundancy. If a server in a load balancing group is unavailable, either due to scheduled maintenance or an unplanned outage, requests to the corresponding domain will be balanced among the remaining servers.

Our Route solution allows the following types of load balancing configurations:

Type	Description
CNAME Record	Load balance traffic that points to a CNAME record in a primary zone hosted on another DNS service provider.
Subdomain Delegation	Load balance traffic that points to a subdomain of a primary zone hosted on another DNS service provider.
Primary Zone (CNAME and Address Records)	Load balance traffic across A, AAAA, or CNAME records in a primary zone hosted on our Route solution.

Reminder: Load balancing can take advantage of server health monitoring via the DNS Health Checks module. This module will incur additional charges on your account. Please contact your DNS account manager for more information.

How Does It Work?

This section describes the distribution of traffic for a load balanced group and how it is affected by a planned or unplanned server outage.

DNS & Load Balancing

A load balancing configuration allows our authoritative DNS servers to distribute requests across various servers or CNAMEs. Five key steps during this process are described below.

1. In response to a client's action (e.g., requesting a web page), an application submits a DNS query to a recursive DNS server.
2. A recursive DNS server forwards the DNS query to the appropriate authoritative name server (i.e., Route). The appropriate authoritative name server is determined via root and top-level domain name servers.
3. A Route name server provides an answer to the DNS query. In this example, it will resolve a domain to an IP address according to your load balancing configuration.
    - **Address Records:** The Route name server will proportionately resolve DNS queries to each address record in the group. The proportion of requests that will be resolved to each address record is defined by its weight.
    - **CNAME Records:** The Route name server will proportionately resolve DNS queries to each CNAME record in the group. The proportion of requests that will be resolved to each CNAME record is defined by its weight. After which, the chosen CNAME record will be resolved to an address record according to its DNS configuration.
4. A recursive DNS server caches the response according to the TTL (e.g., 300 seconds) and forwards it to the client's application.
5. The client's application uses the response to fulfill the request (e.g., direct an HTTP GET request to the site's IP address).

The following illustration depicts how requests are load balanced between three different servers.

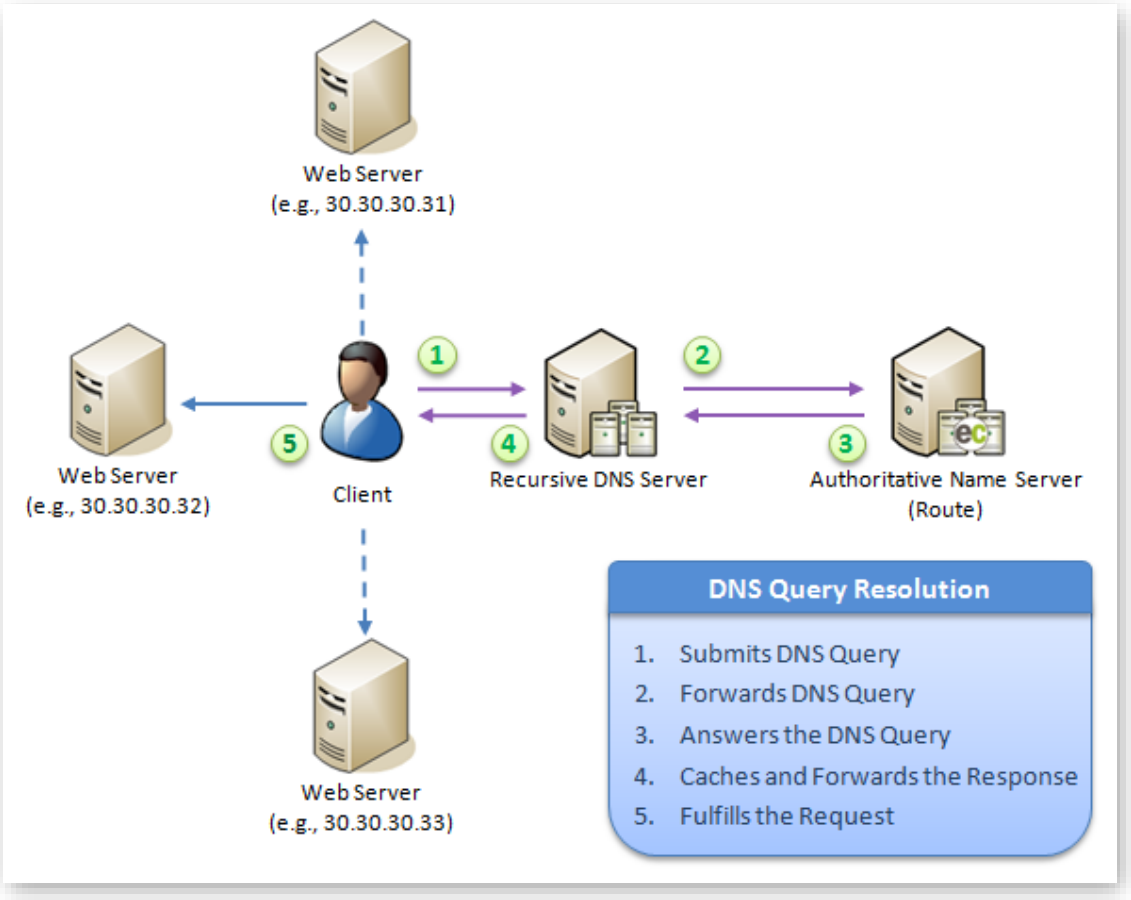

DNS Resolution with Load Balancing

The manner in which requests will be distributed between the above servers is determined by the weight assigned to each server. More heavily weighted servers (i.e., servers that have been assigned a larger value) will receive more requests than servers that have been assigned a lower weight. The following illustration depicts a load balancing configuration that contains three servers. One of the servers has been assigned a lower load balancing weight than the other two. As a result, less traffic is sent to it.

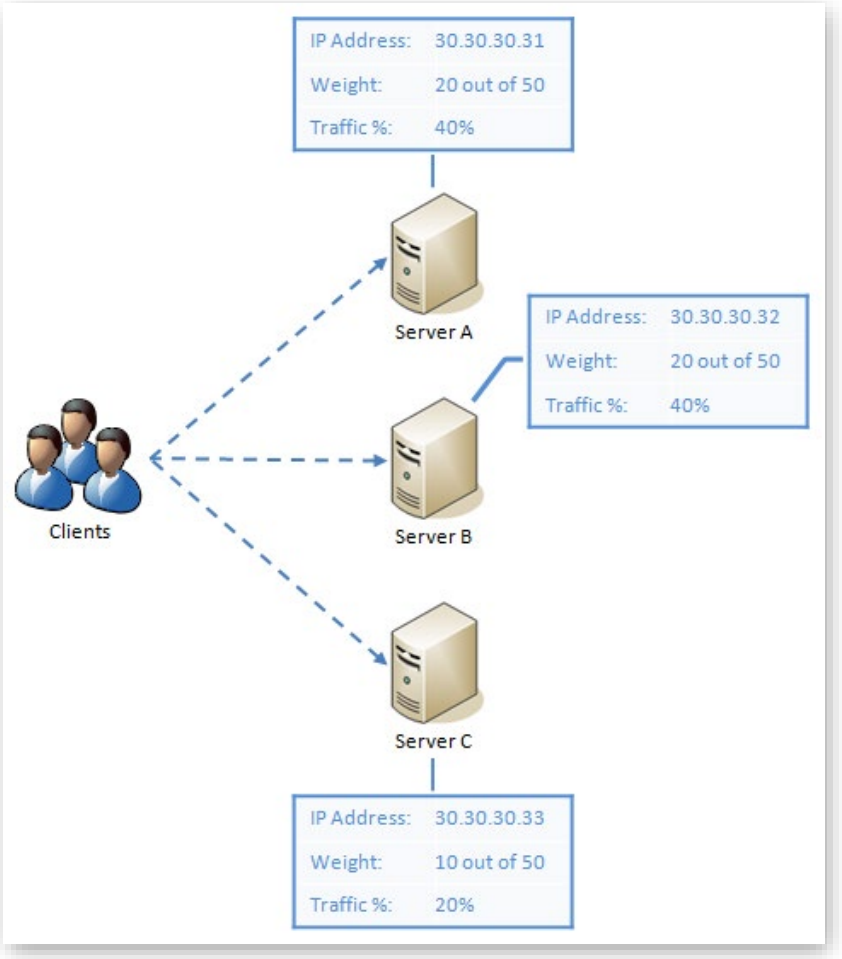

Load Balancing Weight & Request Distribution

It is important to note that proportion of traffic that will be sent to a server/CNAME can be calculated through the following formula:

```
Traffic_Percentage = (Assigned_Weight/Total_Weight) * 100
```

In the above illustration, two servers were assigned a weight of 20 and one server was assigned a weight of 10. The total weight for that load balancing group is 50. This value is calculated by summing the weight assigned to each server (20 + 20 + 10 = 50). The following table uses the above formula to calculate the percentage of traffic that will be sent to each server.

Server	Weight	Calculation	Traffic Percentage
A	20	(20/50) * 100	40%
B	20	(20/50) * 100	40%
C	10	(10/50) * 100	20%
Total:	50	**Total:**	100%

The following illustration demonstrates how traffic is redistributed among the remaining servers when a server is removed from a load balancing configuration.

Note: Our health check system will automatically stop directing traffic to a server when it determines that it is unhealthy as defined by your configuration.

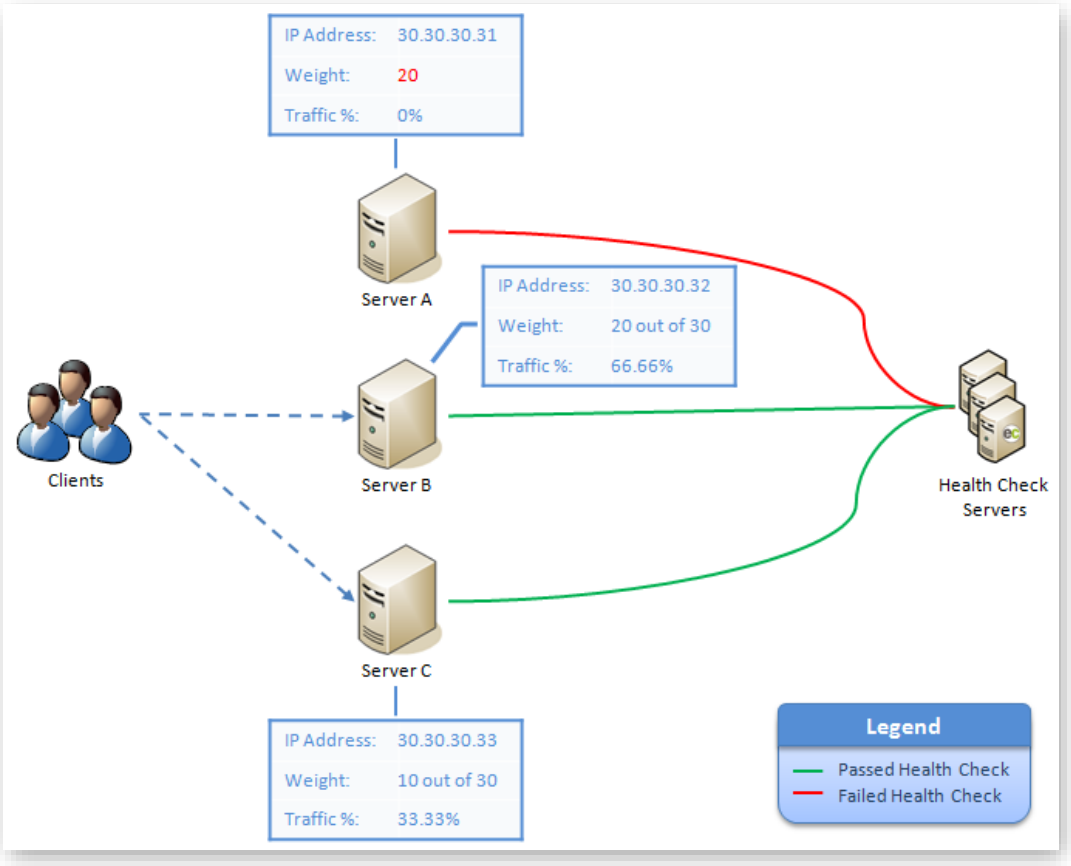

Traffic Redistribution due to the Removal of a Server from Load Balancing

In the above illustration, our DNS service stopped serving traffic to a server that was previously receiving 40% of total traffic. This traffic must now be redistributed proportionally among the remaining two servers. We accomplish this by recalculating the traffic percentage using the new total weight value. The new total weight value, which is 30, can be calculated by summing the weight of all active servers in the configuration (20 + 10). The following table uses the above formula to calculate the percentage of traffic that will be sent to each server.

Server	Weight	Calculation	Traffic Percentage
A	0	(0/30) * 100	0%
B	20	(20/30) * 100	66.67%
C	10	(10/30) * 100	33.33%

Server	Weight	Calculation	Traffic Percentage
Total:	30	**Total:**	100%

In summary, we can make the following conclusions:

- Load balancing provides the means through which traffic can be balanced between servers with varying hardware and/or performance levels.
- A server's weight is important, since it is an indication of the amount of traffic that should be distributed to it.
- Weight is relative to the total weight associated with a load balancing configuration.
    - **Scenario A:** This scenario shows how a server with a weight of 20 will receive the highest proportion of traffic (i.e., 66.66%).

Server	Weight	Calculation	Traffic Percentage
1	5	(5/30) * 100	16.67%
2	5	(5/30) * 100	16.67%
3	20	(20/30) * 100	66.66%
Total:	30	**Total:**	100%

    - **Scenario B:** This scenario shows how the amount of traffic that will be delivered to the same server is dramatically reduced (i.e., 8.89%) when the weight assigned to another server in the group is increased to 200.

Server	Weight	Calculation	Traffic Percentage
1	200	(200/225) * 100	88.89%
2	5	(5/225) * 100	2.22%
3	20	(20/225) * 100	8.89%
Total:	225	**Total:**	100%

- Load balancing allows traffic to be redistributed to healthy servers when a server is pulled out of the load balancing configuration.
- A server's relative weight plays an even greater role when a load bearing server is pulled from a load balancing configuration.
    - **Scenario A:** Notice that server 3, which has a weight of 20, is receiving 22.22% of all traffic.

Server	Weight	Calculation	Traffic Percentage
1	10	(10/90) * 100	11.11%
2	10	(10/90) * 100	11.11%

Server	Weight	Calculation	Traffic Percentage
3	20	(20/90) * 100	22.22%
4	50	(50/90) * 100	55.56%
Total:	90	**Total:**	100%

- **Scenario B:** This scenario shows how traffic is redistributed when our name servers stop serving traffic to server 4. Although all remaining servers are now receiving twice the amount of traffic, the importance of server 3 has risen since 50% of all traffic is now being directed to it.

Server	Weight	Calculation	Traffic Percentage
1	10	(10/40) * 100	25%
2	10	(10/40) * 100	25%
3	20	(20/40) * 100	50%
4	0	(0/40) * 100	0%
Total:	40	**Total:**	100%

Configuration

Keep the following information in mind when configuring a load balancing group:

- The name assigned to a load balancing group must be unique.
- A group can load balance either CNAME or address records, but not to both.
- Although a group can contain a set of IP addresses for each version (i.e., IPv4 and IPv6), traffic is segregated and load balanced according to the IP version used by the network through which the request was received.
    - In a load balancing group that contains both IP versions (i.e., IPv4 and IPv6), a server's weight only affects traffic distribution for servers of the same IP version.
    - A group must contain at least two IP addresses of the same version (i.e., IPv4 or IPv6).
- The speed at which traffic can be completely pulled from an unhealthy server/CNAME is affected by the length of the load balancing group's TTL. TTL configuration varies by module.
    - **Fully Managed:** A load balancing group's TTL must be defined upon creation. Consider setting a short TTL (e.g., 300 seconds) on your load balancing group to ensure that traffic is pulled in a timely manner.
    - **CNAME & Subdomain Delegation:** The TTL for all load balancing groups is 300 seconds. This setting cannot be modified.

- A health check configuration should be added to each IP address/CNAME to ensure that traffic is redistributed when a server can no longer serve traffic.
- A load balancing group will not be created, modified, or deleted until **Submit Group** is clicked.

Load Balancing Group Creation

In order to load balance traffic, it must flow through our DNS servers. However, this does not mean that your zone must be fully managed by our DNS service. A brief description is provided below on the various sources through which traffic can be routed to our DNS service for the purpose of load balancing.

Type	Description
Primary Zone	The traffic flowing through a zone that is managed by our DNS service can easily be load balanced. Simply include the following in a zone's load balancing group: • Set of address records (A or AAAA). • Set of CNAME records **Note:** The hostname and TTL values associated with the selected records must be consistent across a load balancing group. Therefore, upon saving a load balancing configuration, the hostname and TTL settings associated with the corresponding records will be overwritten.
CNAME Record	All traffic that flows through a CNAME record can be load balanced to either address records or even other CNAME records. This configuration requires pointing a CNAME record, the one to which requests will be directed, to the domain associated with a load balancing configuration (e.g., mycname.0001.edgecastdns.net). **Warning:** This type of configuration is dependent on a DNS service provider's ability to resolve a CNAME record to our hostname. We will be unable to provide traffic management services if another DNS service provider experiences a service outage. **Note:** A CNAME record can be created from the DNS service provider that owns the desired zone.

Type	Description
Subdomain Delegation	A subdomain can be delegated to our DNS service. Traffic directed to that subdomain can be load balanced across address or CNAME records. Keep the following information in mind when delegating a subdomain: • From your DNS service provider, create NS records that point to our authoritative name servers (i.e., ns1.edgecastdns.net, ns2.edgecastdns.net, ns3.edgecastdns.net, and ns4.edgecastdns.net). • The name assigned to the load balancing group should match the domain associated with the above NS records (e.g., www.mydomain.com). **Warning:** This type of configuration is dependent on a DNS service provider's ability to indicate that our name servers can provide an authoritative answer for the desired subdomain. We will be unable to provide traffic management services if another DNS service provider experiences a service outage. **Note:** This option is typically used to hide the DNS service provider (i.e., Edgecast) through which a domain will be resolved.

The procedure through which a load balancing group can be created varies by the method used to direct traffic to DNS service. Instructions for each method are provided below.

To create a load balancing group (Primary Zone)

1. If you haven't already created the desired zone, then please do so now. Information on how to create a zone can be found in the **DNS Zone Management** chapter.
2. Make sure that the desired records (i.e., A, AAAA, or CNAME) have been created in that zone.
3. Click **Create LB/FO**.
4. Click **Create Load Balancing**.
5. Mark the records that will be load balanced. Make sure to only select records of the same type (i.e., address or CNAME).
6. Click **Finish creating LB**. A dialog box will appear.
7. In the **Name** option, define the hostname for the load balancing group.
8. In the **TTL** option, define the number of seconds that the records should remain cached.
9. Click **Add**.

10. If you would like to add a health check or modify the weight assigned to a record, perform the following steps:
    i. Click **Edit** next to the desired IP address/CNAME.
    ii. Under the **Weight** column, define the load balancing weight for that item.
    iii. The **Add HC** button provides the means through which a health check configuration can be created. For detailed information on how to create a health check configuration, please refer to the **Health Checks** chapter.
11. Repeat the previous step as needed.
12. When finished, click **Done**.
13. Click **Submit Group** to save your changes.

Important: Creating a load balancing group will overwrite the name and TTL assigned to the records associated with it. If this is undesired, you may need to create additional records. Keep in mind that you can create multiple records that point to the same IP address/domain.

To create a load balancing group (CNAME Record/Subdomain Delegation)

1. Navigate to the **Route (DNS)** page.
2. Click **Add New**.
3. In the **Type** option, select "Load Balancing."
4. In the **Name** option, type the desired name (e.g., www). Additional information is provided below.
    - **CNAME Record:** A system-defined domain (e.g., www.0001.edgecastdns.net) will be generated from the specified hostname.
    - **Subdomain Delegation:** The name assigned to the load balancing group should match the subdomain delegation (e.g., www or ftp).
5. In the **Type** option, select whether traffic will be directed to our DNS service via a CNAME record or subdomain delegation.
6. Click **Add**.
7. Add either an IP address or a CNAME to the load balancing group.
    i. Click **Add**.
    ii. In the **Value** option, specify either an IPv4/IPv6 address or a CNAME's domain.
    iii. In the **Type** option, select whether the item defined in the **Value** option is an IPv4/IPv6 address or a CNAME.
    iv. In the **Weight** option, define the load balancing weight that will be assigned to this IP address.

v. Click **Add**.

8. Repeat step 7 for each IP address or CNAME that should be included in the load balancing group. Keep in mind that the **Type** option is only defined once.
9. If you would like to add a health check, perform the following steps:
    i. Click **Edit** next to the desired IP address or CNAME.
    ii. The **Add HC** button provides the means through which a health check configuration can be created. For detailed information on how to create a health check configuration, please refer to the **Health Checks** chapter.
    iii. Click **Verify and Save** when finished.
    iv. Repeat steps i – iii for each desired IP address.
10. Click **Submit Group** to save your changes.
11. Perform one of the following:
    - **CNAME Record:** Create a CNAME record in a primary zone hosted on another DNS service provider that points to the system-defined domain (e.g., www.0001.edgecastdns.net) generated for the load balancing group. This domain can be viewed by clicking  to expand the Group Information section.
    - **Subdomain Delegation:** From your alternate DNS service provider, delegate the subdomain (e.g., www.mydomain.com) to Route name servers. Our Route name servers can be viewed by clicking  to expand the Group Information section.

Load Balancing Group Modification

There are a variety of reasons for updating a load balancing group, such as:

- Add or remove servers or CNAMEs from the group.
- Change the load balancing weight assigned to a server or a CNAME.
- Add, modify, or delete the health check configuration associated with an IP address or a CNAME. For detailed information on how to administer a health check configuration, please refer to the **Health Checks** chapter.
- Reintegrate a server/CNAME back into a load balancing group. Our DNS service will stop serving traffic to a server/CNAME that has failed a majority of health checks. The health check configuration defined for that server or CNAME determines whether manual intervention is required to indicate when our Route name servers can resume serving traffic to it.

Information on how to perform common tasks is provided below.

To add a server or CNAME to a load balancing group (Primary Zone)

1. From the **Route (DNS)** page, click on the desired zone.
2. Click **Edit** next to the desired load balancing group.
3. Add the desired server or CNAME.
    i. Click **Add**.
    ii. Perform one of the following:
        - **CNAME:** In the **CNAME** option, specify a domain to which traffic will be load balanced.
        - **Server:** In the **IP Address** option, specify the IP address to which traffic will be load balanced.
    iii. In the **Weight** option, define the load balancing weight that will be assigned to it.
    iv. Click **Add**.
4. Click **Done**.
5. Click **Submit Group** to save your changes.

To add a server or CNAME to a load balancing group (CNAME record & Subdomain delegation)

1. From the **Route (DNS)** page, click on the desired load balancing group.
2. Click **Add**.
3. In the **Value** option, specify the IP address or the domain to which traffic will be load balanced.
4. In the **Weight** option, define the load balancing weight that will be assigned to the IP address or domain.
5. Click **Add**.
6. Click **Submit Group** to save your changes.

To remove a server or CNAME from a load balancing group

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Zone:** Click on the desired zone and then click **Edit** next to the desired load balancing group.
    - **CNAME record & Subdomain Delegation:** Click on the desired load balancing group.
2. Click **Edit** next to the desired IP address or domain.
3. Click  .
4. **Zone Only:** Click **Done**.
5. Click **Submit Group** to save your changes.

To change the assigned load balancing weight

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Zone:** Click on the desired zone and then click **Edit** next to the desired load balancing group.
    - **CNAME record & Subdomain Delegation:** Click on the desired load balancing group.
2. Click **Edit** next to the desired IP address or domain.
3. Set the desired weight under the **Weight** column.
4. Click **Done**.
5. Click **Submit Group** to save your changes.

To manually reintegrate a server or CNAME

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Zone:** Click on the desired zone and then click **Edit** next to the desired load balancing group.
    - **CNAME record & Subdomain Delegation:** Click on the desired load balancing group.
2. Click **Edit** next to the desired IP address or domain.
3. Click **Edit** under the **Health Check** column.
4. Click **Manually Reintegrate**.
5. Click **Verify and Save**.
6. Click **Done**.
7. Click **Submit Group** to save your changes.

Load Balancing Group Deletion

A load balancing group can be deleted at any time by performing one of the following procedures.

To delete a load balancing group (Primary Zone)

1. From the **Route (DNS)** page, click on the desired zone.
2. Click **Edit** next to the desired load balancing group.
3. Click **Edit** next to the desired IP address or domain.
4. Click  .
5. Repeat steps 3 and 4 until the load balancing group no longer contains records.
6. Click **Done**.
7. Click **Submit Group** to save your changes.

To delete a load balancing group (CNAME & Subdomain Delegations)

1. Navigate to the **Route (DNS)** page.
2. Hover the cursor over the desired load balancing group.
3. Click the delete icon () that appears next to it.
4. When prompted, confirm the deletion of that load balancing group.

DNS Failover

Overview

A primary/backup relationship can be established between two servers or domains. This type of relationship allows a backup server/domain to take over when the primary server/domain can no longer fulfill its responsibilities. This prevents an outage from impacting site traffic. This process is known as failover.

Our Route solution allows the following types of failover configurations:

Type	Description
CNAME Record	Failover traffic that points to a CNAME record in a primary zone hosted on another DNS service provider.
Subdomain Delegation	Failover traffic that points to a subdomain of a primary zone hosted on another DNS service provider.
Primary Zone (Address Records)	Failover traffic across multiple address or CNAME records in a primary zone hosted on our Route solution.

Reminder: Failover requires server health monitoring via the DNS Health Checks module. This module will incur additional charges on your account. Please contact your DNS account manager for more information.

How Does It Work?

A failover configuration establishes a primary and backup relationship between two servers or domains. Our authoritative name servers will send all traffic to the primary server/domain until it fails a majority of its health checks. At which point, all traffic will be redirected to the backup server/domain.

The following illustration shows the process through which all requests are resolved to the IP address of a primary server.

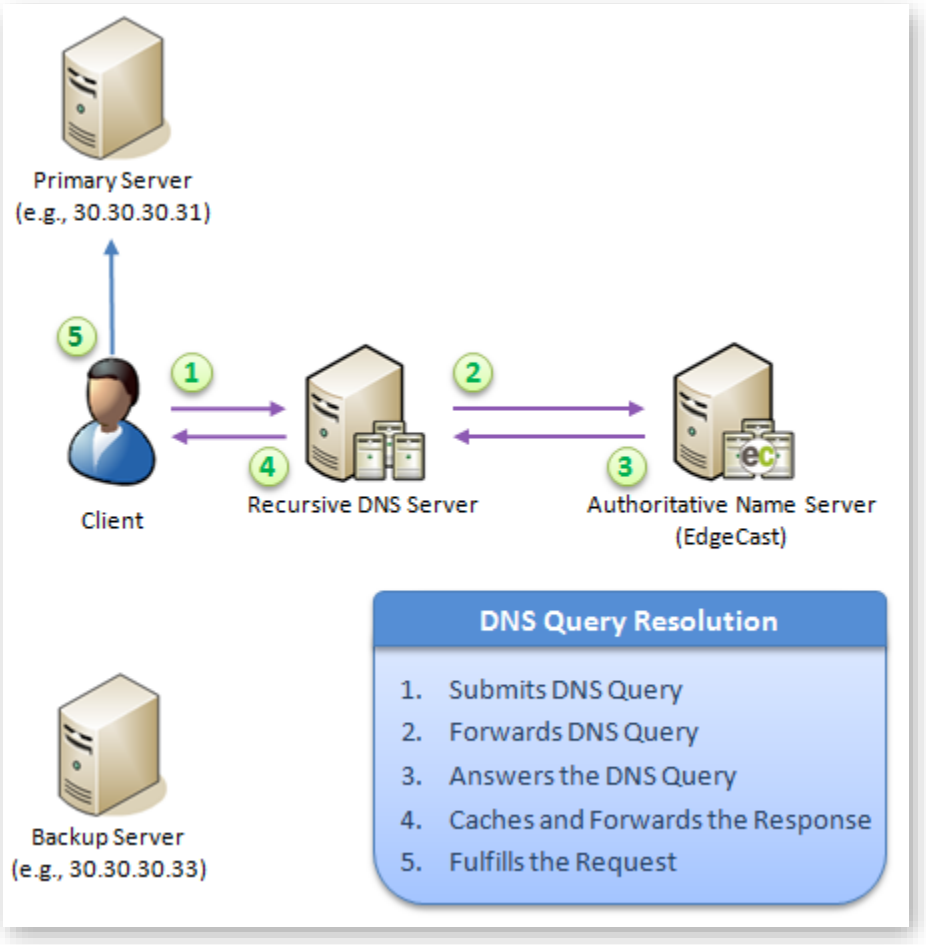

Resolving a DNS Query for a Failover Configuration

A health check configuration plays a major role in determining when to fail traffic over to a backup server/domain. Once a health check configuration has been established, a server/domain will be polled at regular intervals to ensure that it is still healthy.

In the following illustration, both the primary and backup servers have passed a majority of their health checks. All traffic is directed to the primary web server as a result of a failover configuration.

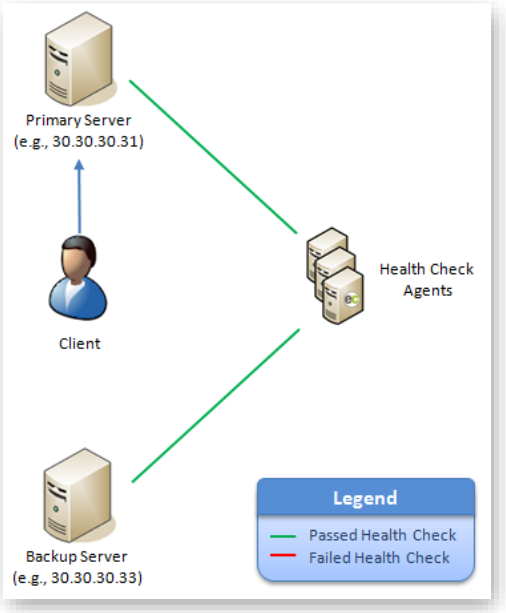

Healthy Primary Server

If the primary server fails a majority of its health checks, our DNS service will redirect all traffic to the backup server. This is illustrated below.

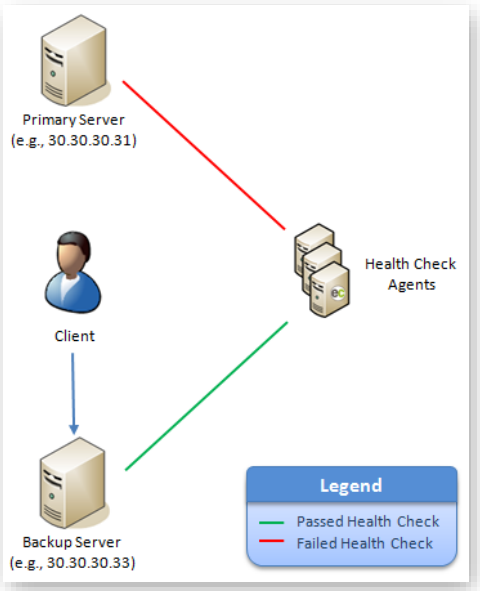

Unhealthy Primary Server

In the above scenario, traffic will continue to be served to the backup server until the primary server is reintegrated back into the failover configuration. At which point, our DNS service will redirect all traffic to the primary server.

Configuration

Keep the following information in mind when configuring a failover group:

- The name assigned to a failover group must be unique.
- A group can failover either CNAME or address records, but not to both.
- Although a group can contain a set of IP addresses for each version (i.e., IPv4 and IPv6), traffic is segregated according to the IP version used by the network through which the request was received.
    - In a failover group that contains both IP versions (i.e., IPv4 and IPv6), the role (i.e., primary or backup) assigned to a server is specific to that IP version.
    - Only a single IP address per IP version can be designated as the primary server. Setting the **Primary option** will toggle the status of the other IP address in the same group.
    - A group must contain at least two IP addresses of the same version (i.e., IPv4 or IPv6).
- A health check configuration should be added to each IP address/CNAME to ensure that traffic is redirected when a server/domain can no longer properly serve traffic.
- The speed at which traffic can be completely pulled from an unhealthy server/domain is affected by the length of the failover group's TTL. TTL configuration varies according to the type of DNS configuration being performed.
    - **Primary Zone:** A failover group's TTL must be defined upon creation. Consider setting a short TTL (e.g., 300 seconds) on a failover group to ensure that traffic is pulled in a timely manner.
    - **CNAME & Subdomain Delegation:** The TTL for all failover groups is 300 seconds. This setting cannot be modified.
- A failover group will not be created, modified, or deleted until **Submit Group** is clicked.

Failover Group Creation

A brief description is provided below on the various sources through which traffic can be routed to our DNS service for the purpose of failover.

Type	Description
Zone	A failover configuration can be easily established for the traffic flowing through a zone that is managed by our DNS service. It is simply a matter of defining the address or CNAME records that will be part of the failover configuration. **Note:** The hostname and TTL values associated with address records must be consistent across a failover group. Therefore, upon saving a failover configuration, the hostname and TTL settings associated with the corresponding records will be overwritten.
CNAME Record	A failover configuration can be defined for a CNAME record. This type of configuration defines a primary and a backup IP address/domain to which that CNAME record will be resolved. A failover group requires that the desired CNAME record point to the domain associated with a failover configuration. This step must be performed from the DNS service provider that owns the zone. **Warning:** This type of configuration is dependent on the DNS service provider's ability to resolve a CNAME record to our hostname. We will be unable to provide traffic management services if another DNS service provider experiences a service outage.
Subdomain Delegation	A subdomain can be delegated to our DNS service. Keep the following information in mind when delegating a subdomain. • From your DNS service provider, create NS records that point to our authoritative name servers (i.e., ns1.edgecastdns.net, ns2.edgecastdns.net, ns3.edgecastdns.net, and ns4.edgecastdns.net). • The name assigned to the failover group should match the domain associated with the above NS records (e.g., www.mydomain.com). **Warning:** This type of configuration is dependent on the DNS service provider's ability to indicate that our name servers can provide an authoritative answer for your subdomain. We will be unable to provide traffic management services if another DNS service provider experiences a service outage. **Note:** This option is typically used to hide the DNS service provider (i.e., Edgecast) through which a domain will be resolved.

The procedure through which a failover group can be created varies by the method used to direct DNS traffic. Instructions for each method are provided below.

To create a failover group (Primary Zone)

1. Make sure that the desired zone has already been created. Information on how to create a zone can be found in the **DNS Zone Management** chapter.
2. Make sure that the desired address or CNAME records have been created in that zone.
3. Click **Create LB/FO**.
4. Click **Create Failover**.
5. Mark either two address or CNAME records.
6. Click **Finish creating FO**. A dialog box will appear.
7. In the **Name** option, define the subdomain for the failover group.
8. In the **TTL** option, define the number of seconds that the records associated with this group should remain cached.
9. Click **Add**.
10. Click **Edit** corresponding to the IP address or CNAME that should be the primary server.
11. Make sure that the check box in the **Primary** column is marked.
12. The **Add HC** button provides the means through which a health check configuration can be created. For detailed information on how to create a health check configuration, please refer to the **Health Checks** chapter.
13. Click **Edit** corresponding to the IP address or CNAME that should be the backup server.
14. Make sure that the check box in the **Primary** column is cleared.
15. Add a health check configuration for the backup server/domain.
16. Click **Done**.
17. Click **Submit Group** to save your changes.

Important: Creating a failover group will overwrite the name and TTL assigned to the records associated with it. If this is undesired, you may need to create additional records. Keep in mind that you can create multiple records that point to the same IP address/domain.

To create a failover group (CNAME Record/Subdomain Delegation)

1. Navigate to the **Route (DNS)** page.
2. Click **Add New**.
3. In the **Type** option, select "Failover."
4. In the **Name** option, type one of the following:
    - **CNAME Record:** A system-defined domain (e.g., www.0001.edgecastdns.net) will be generated from the specified hostname.
    - **Subdomain Delegation:** The name assigned to the failover group should match the subdomain delegation (e.g., www or ftp).
5. In the **Type** option at the bottom of the dialog box, select whether traffic will be directed to our DNS service via a CNAME record or subdomain delegation.
6. Click **Add**.
7. Add either an IP address or a CNAME to the failover group.
    i. Click **Add**.
    ii. In the **Value** option, specify either an IP address or a CNAME's domain.
    iii. In the **Type** option, select whether the failover group will contain IP addresses or domains (CNAME).
    iv. In the **Primary** option, indicate whether the current configuration is for a primary server/domain.
    v. Click **Add**.
8. Repeat step 7 for each IP address or CNAME that should be included in the failover group. Keep in mind that the **Type** option is only defined once.
9. Create a health check configuration for each IP address/domain associated with your failover configuration.
    i. Click **Edit** next to the desired IP address.
    ii. The **Add HC** button provides the means through which a health check configuration can be created. For detailed information on how to create a health check configuration, please refer to the **Health Checks** chapter.
    iii. Click **Done** when finished.
    iv. Repeat steps i – iii for the second IP address/domain.
10. Click **Submit Group** to save your changes.

11. Perform one of the following:
    - **CNAME Record:** Create a CNAME record in a primary zone hosted on another DNS service provider that points to the system-defined domain (e.g., www.0001.edgecastdns.net) generated for the failover group. This domain can be viewed by clicking  to expand the **Group Information** section.
    - **Subdomain Delegation:** From your alternate DNS service provider, delegate the subdomain (e.g., www.mydomain.com) to Route name servers. Our Route name servers can be viewed by clicking  to expand the **Group Information** section.

Failover Group Modification

Common reasons for updating a failover group are to:

- Modify an IP address or domain. This is the quickest method to point a failover group to a different server/domain.
- Designate a different primary server/domain. Setting up a different primary server/domain will redirect all traffic it.
- Remove a server/domain from a failover group. Keep in mind that a failover group must have exactly two:
    - IP addresses of the same type (i.e., IPv4 or IPv6)
    - CNAMEs
- Add, modify, or delete the health check configuration associated with an IP address/domain. For detailed information on how to administer a health check configuration, please refer to the **Health Checks** chapter.
- Reintegrate a server/CNAME back into a load balancing group. Our DNS service will not serve traffic to a server/CNAME that has failed a majority of health checks. The health check configuration defined for that server or CNAME determines whether manual intervention is required to indicate when our Route name servers can resume serving traffic to it.

Information on how to perform common tasks is provided below.

To modify server/domain properties

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Primary Zone:** Click on the desired zone and then click **Edit** next to the desired failover group.
    - **CNAME Record & Subdomain Delegation:** Click on the desired failover group.
2. Click **Edit** next to the desired IP address or domain.
3. Update the assigned IP address or domain.
4. Verify its primary/backup status under the **Primary** column.
5. Add or modify its health check configuration.
6. Click **Done**.
7. Click **Submit Group** to save your changes.

To remove a server from a failover group

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Primary Zone:** Click on the desired zone and then click **Edit** next to the desired failover group.
    - **CNAME Record & Subdomain Delegation:** Click on the desired failover group.
2. Click **Edit** next to the desired IP address or domain.
3. Click  .
4. **Zone Only:** Click **Done**.
5. Click **Submit Group** to save your changes.

To manually reintegrate a server

1. Perform one of the following procedures from the **Route (DNS)** page:
    - **Primary Zone:** Click on the desired zone and then click **Edit** next to the desired failover group.
    - **CNAME Record & Subdomain Delegation:** Click on the desired failover group.
2. Click **Edit** next to the desired IP address or domain.
3. Click **Edit** under the **Health Check** column.
4. Click **Manually Reintegrate**.
5. Click **Verify and Save**.
6. Click **Submit Group** to save your changes.

Failover Group Deletion

A failover group can be deleted at any time by performing one of the following procedures.

To delete a failover group (Primary Zone)

1. From the **Route (DNS)** page, click on the desired zone.
2. Click **Edit** next to the desired failover group.
3. Click **Edit** next to the desired IP address or domain.
4. Click  .
5. Repeat steps 3 and 4.
6. Click **Done**.
7. Click **Submit Group** to save your changes.

To delete a failover group (CNAME & Subdomain Delegations)

1. Navigate to the Route (DNS) page.
2. Hover the cursor over the desired failover group.
3. Click the delete icon () that appears next to it.
4. When prompted, confirm the deletion of that failover group.

Secondary DNS

Overview

Secondary DNS allows our name servers to be authoritative for DNS zones managed outside of our network.

This type of setup allows the following:

- Automatic propagation of updates from DNS zones managed on a master name server. This allows your zone content to be managed elsewhere, while still leveraging our global Anycast DNS product and network to resolve your clients' DNS queries.
- Hidden master name servers. This security measure protects your name server from malicious attacks while continuing to manage your zones locally or through another DNS service provider. This allows you to leverage Route's performance, reliability, and security capabilities while maintaining continuity with your existing zone management process.

How Does It Work?

The purpose of our Secondary DNS solution is to serve externally managed zones through our service. In order to achieve this goal, our Secondary DNS solution synchronizes zone content between a master name server and our Anycast DNS system. This process consists of the following steps:

1. Initial Zone Transfer
2. Zone Synchronization

This section will explain both of these steps and the manner in which DNS queries are handled by our solution.

Initial Zone Transfer

The following diagram provides an overview of the process through which a zone is transferred to our Route solution.

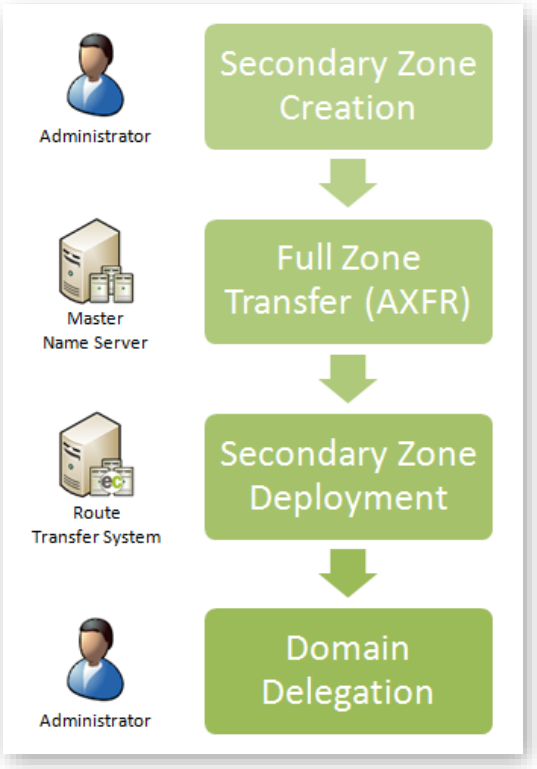

Initial Zone Transfer & Domain Delegation

The following steps describe this process in more detail.

1. An administrator creates a secondary zone within Route.
2. The Route transfer system requests a zone transfer for that zone from a master name server.
3. The master name server performs a full zone transfer (AXFR) for the requested secondary zone.
4. A read-only copy of that zone is created within Route. Glue records for our vanity name servers are automatically added to that zone. Vanity name servers lend a professional appearance to your site's DNS. The resulting zone is known as a secondary zone.
5. Our Route transfer system will then deploy the secondary zone to our Route name servers.

6. DNS queries must now be directed to our Route name servers. This step requires an administrator to update the appropriate domain registrar to point the zone to our vanity name servers.

Zone Synchronization

Once a secondary zone has been created, our DNS solution ensures that it remains synchronized with the master zone through the following workflow:

1. **Serial Number Check:** Our DNS solution checks whether the master zone has been updated approximately every 120 seconds. This check consists of the following steps:
    i. Our system will request the zone's SOA record from each master name server in the master server group associated with the secondary group.
    ii. The secondary zone's serial number will be compared against the highest serial number returned in step 1.
2. **Full Zone Transfer (AXFR):** If the master zone's serial number has been incremented, then our DNS solution will submit an AXFR query to the corresponding master name server. The secondary zone will be replaced with the full zone returned by the master name server.

Tip: View the raw data by loading the secondary zone and then clicking the "Transfer Data" link corresponding to the desired master name server. Master name servers are listed in the **Master Server Group** section.

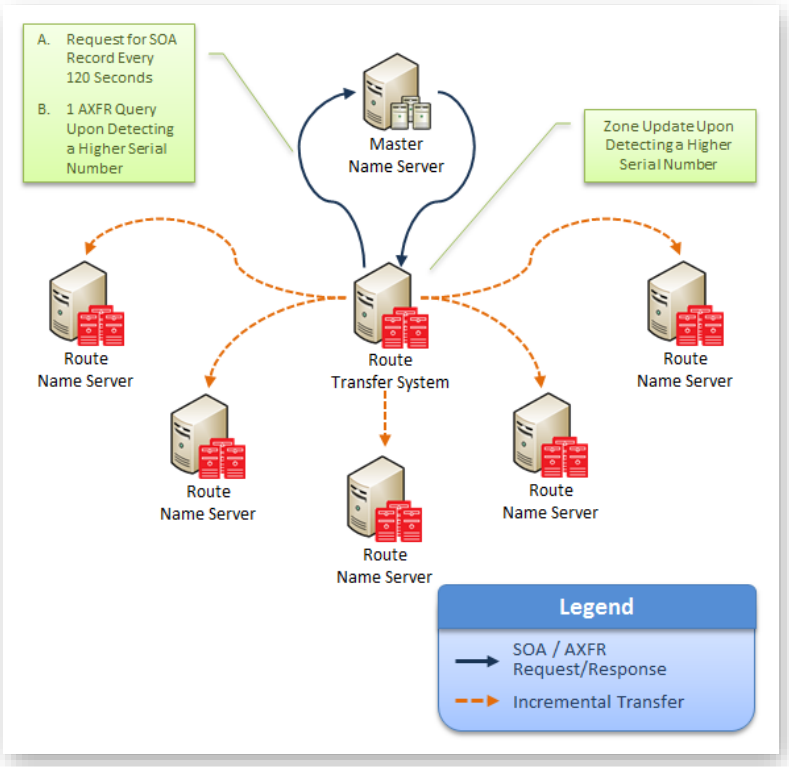

Secondary Zone Synchronization

DNS Query Handling

Our Route name servers will be able to authoritatively resolve DNS queries to your zone once the following conditions are met:

- The secondary zone in question has been deployed to our Route name servers.
- The zone has been delegated to our vanity name servers.

If your master name servers are unavailable, our Route solution will continue to serve the last transferred zone regardless of whether the requested record's TTL has expired. This differentiates our service from traditional Master/Slave DNS systems where slave name servers stop serving zones when the master name server remains unavailable beyond a record's TTL.

This design ensures that stale DNS requests are resolved efficiently even when your master name server is unavailable. In the following illustration, notice how DNS requests handled by our name servers are never forwarded to your master name server.

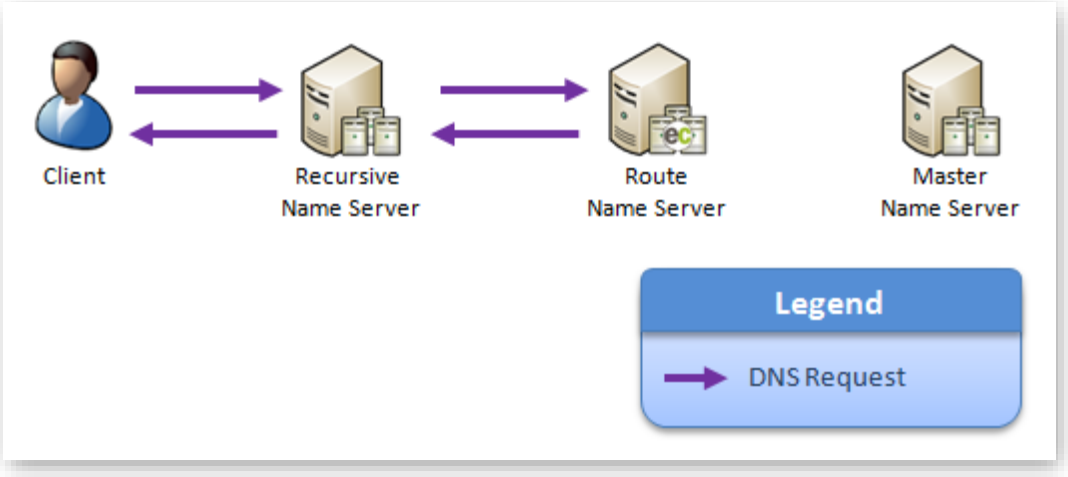

Route Name Servers Behave as Master Name Servers

Configuration

Secondary DNS configuration requires the creation of the following items:

Item	Description	Task(s)
Master Server Group	Allows the grouping of master name servers. A master server group defines the set of name servers from which zone data can be transferred.	Create a master server group and add at least one master server to it.
Secondary Zone Group	Defines the set of secondary zones that will be imported from the name servers that make up a master server group. **Note:** A secondary zone may only be modified or deleted from its corresponding secondary zone group.	Create a secondary zone group and then perform the following steps: 1. Assign a master server group to it. 2. If needed, add/assign TSIG keys. 3. Define each zone that will be transferred to our Route name servers.

Upon saving a secondary zone group, a full zone transfer (AXFR) will be performed for each new zone. Each secondary zone will contain the following:

- All of the records, including DNSSEC records, associated with the original zone.
- Required NS and Host records for our vanity name servers.

Note: Once you are ready to start serving production traffic through a secondary zone, you should add our vanity name servers as glue records in your domain's registrar. For more information, please refer to the **Switching DNS Service Provider** section below.

Note: The naming convention for our vanity name servers is s#ns#.*zonename* (e.g., s1ns1.myzone.com).

Note: The records used to define our vanity name servers are segregated from other zone records. These records can be viewed by expanding the **Default Records** section.

Key information:

- It is possible that a master zone may contain certain types of unsupported records that will not be imported into our system. For this reason, we provide raw transfer data that you may compare against the records imported into our system. We also provide information on any transmission errors that may have occurred during a zone transfer.
- All secondary zone records are read-only. Our secondary zones are automatically updated to reflect changes to the source zone.
- Each secondary zone indicates the serial number, master name server, TSIG, and the date/time stamp of when it was transferred to our system.
- Make sure to authorize zone transfer requests from our Route transfer system to master name servers.

Master Server Group Administration

A master server group allows you to group and manage master name servers from a single location instead of from each secondary zone group or secondary zone.

Keep in mind the following information:

- The name assigned to a master server group or a master name server is simply a label to allow easy identification of that group or name server.
- A master name server can belong to multiple master server groups. If you have previously created a master name server in another master server group, you may simply select it when defining the membership of a new master server group.
- TSIG administration can be performed when adding or modifying a secondary zone group. Changes to a TSIG key are immediately reflected across all secondary zone groups. This allows you to make bulk changes to TSIG keys that affect many secondary zone groups and associated secondary zones.
- Once a master server group has been associated with a secondary zone group, its name servers will be queried every 120 seconds for each zone identified in that secondary zone group. For this reason, we recommend that a master server group should only consist of name servers that are authoritative for the same set of zones. Likewise, a secondary zone group should be limited to those zones.

To create a master server group

1. Navigate to the **Route (DNS)** page.
2. Click **Add New**. The **Add New** dialog box will appear.
3. From the **Type** option, select "Master Server Group."
4. In the **Name** option, type the name of the master server group.
5. Click **Add**.
6. Perform the following steps to add a master name server:
    i. Click **Add Master**. The **Add New Master Server** dialog box will appear.
    ii. In the **Name** option, type the name through which the master name server will be identified.
    iii. In the **IP Address** option, type the IPv4 or IPv6 address corresponding to the desired master name server.
    iv. Click **Add**.
7. Click **Submit Group** to save your changes.

To modify a master server group

1. Navigate to the **Route (DNS)** page.
2. Click on the master server group that you would like to modify.
3. Perform the following steps to add a master name server:
    i. Click **Add Master**. The **Add New Master Server** dialog box will appear.
    ii. In the **Name** option, type the name through which the master name server will be identified.
    iii. In the **IP Address** option, type the IPv4 or IPv6 address corresponding to the desired master name server.
    iv. Click **Add**.
4. Perform the following steps to modify a master name server:
    i. Click **Edit** next to the desired master name server.
    ii. Review and/or modify the name and IP address associated with that master name server.
    iii. Click **Done**.
5. Perform the following steps to delete a master name server:
    i. Click **Edit** next to the desired master name server.

ii. Click  .

6. Click **Submit Group** to save your changes.

To delete a master server group

1. Navigate to the **Route (DNS)** page.
2. Hover the cursor over the desired master server group.
3. Click the delete icon () that appears next to it.
4. When prompted, confirm the deletion of that master server group.

Secondary Zone Group Administration

The purpose of a secondary zone group is to define one or more secondary zones. A secondary zone will be created for each zone defined in a secondary zone group. These secondary zones will be populated with:

- A default set of records that identify our vanity name servers.
- The set of records associated with the original zone. These records are retrieved via a full zone transfer (AXFR).

Although the administration of these secondary zones is performed from the secondary zone group, you may only view the records associated with that secondary zone by clicking on it from the **Route (DNS)** page.

Key information:

- A master server group must be assigned to a secondary zone group. All master name servers associated with that group will be queried every 120 seconds per zone to check for updates.
- Administer transaction signature (TSIG) keys when defining a secondary zone group. For more information, please refer to the **Transaction Signature (TSIG) Authentication** section below.
- Each zone must be added individually to a secondary zone group.
- A full zone transfer will take place for each new zone upon saving the secondary zone group.
- A zone should only be added to a single secondary zone group. Duplicate secondary zones are not allowed.
- If your zone name does not end in a period, then one will be appended for you. This is standard DNS notation and indicates that the zone name is fully qualified.
- Deleting a secondary zone group will also delete all of its secondary zones.

To create a secondary zone group

1. Navigate to the **Route (DNS)** page.
2. Click **Add New**. The **Add New** dialog box will appear.
3. From the **Type** option, select "Secondary Zone Group."
4. In the **Name** option, type the name of the secondary zone group.
5. Click **Add**.
6. Select a master server group from the **–Select Master Server Group –** option. The selected group should only contain name servers that are authoritative for the zones that will be defined in this new secondary zone group.
7. Optional. If you would like to authenticate DNS updates through TSIG, then you will need to perform the following steps for each desired master name server:
    i. Define a TSIG key. For more information, please see the **Transaction Signature (TSIG) Authentication** section below.
    ii. Assign a TSIG key to a master name server. A TSIG key can be assigned to a master name server by simply selecting it from the **–Select TSIG –** option that appears next to it.
8. Perform the follow steps for each zone that will be associated with this group:
    i. In the **Add zone** option, type the name of the zone that will be transferred to our DNS service. A period will be automatically appended to zone names that do not end in a period.
    ii. Click **Add**.
9. Click **Submit Group**. A full zone transfer will be performed for each zone associated with this new secondary zone group. These secondary zones can be viewed from the **Route (DNS)** page.

To modify a secondary zone group

1. Navigate to the **Route (DNS)** page.
2. Click on the secondary zone group that you would like to modify.
3. Assign a different master server group by simply selecting it from the **–Select Master Server Group –** option. Keep in mind that the selected master server group should only contain name servers that are authoritative for the zones defined in this secondary zone group.
4. Optional. If you would like to modify TSIG authentication settings, then you will need to perform the following steps for each desired master name server:
    - If DNS updates for a particular master name server should not leverage TSIG authentication, then set the **–Select TSIG –** option to "- Select TSIG -."
    - If you would like to change the TSIG key assigned to a master name server, then simply select the desired key from the **–Select TSIG –** option that appears next to it. If the desired key hasn't been registered, then you should add it. For more information, please see the **Transaction Signature (TSIG) Authentication** section below.
5. Perform the following steps to add a zone:
    i. In the **Add zone** option, type the name of the zone that will be transferred to our DNS service. A period will be automatically appended to zone names that do not end in a period.
    ii. Click **Add**.
6. Perform the following steps to modify a zone:
    i. Click **Edit** next to the desired zone.
    ii. Review and/or modify the name associated with that zone.
    iii. Click **Done**.
7. Perform the following steps to delete a zone:
    i. Click **Edit** next to the desired zone.
    ii. Click  .
8. Click **Submit Group** to save your changes. The following actions will take place upon saving your changes:
    - A full zone transfer will take place for newly added or modified zone. A secondary zone will be added for each transferred zone. If you modified a zone, then the original secondary zone will be deleted.

- If you deleted a zone from the group, then the corresponding secondary zone will be deleted.

To delete a secondary zone group

1. Navigate to the **Route (DNS)** page.
2. Hover the cursor over the desired secondary zone group.
3. Click the delete icon () that appears next to it.
4. When prompted, confirm the deletion of that secondary zone group.

Transaction Signature (TSIG) Authentication

Transaction Signature (TSIG) allows our name servers to authenticate communication to your master name server(s). This requires that you register a TSIG key and then assign it to a master name server. A brief description for each component of a TSIG key is provided below.

Component	Description
TSIG Alias	A label that provides a brief description for a TSIG key. This label is solely used to identify a TSIG key when assigning it to a master name server.
Key Name	Identifies the key on the master name server and our Route name servers. This name must be unique.
Key Type	Identifies the cryptographic hash function used to generate the key value. Our TSIG implementation supports MD5, SHA-1, and SHA-2 (SHA-224, SHA-256, SHA-384, and SHA-512).
Key Value	Identifies a hash value through which our name servers will be authenticated to a master name server.

Key information:

- A TSIG key may be assigned to each master name server associated with a secondary zone group. Our name servers will use the selected key when communicating with those master name servers.
- If you plan on using TSIG authentication, it is recommended to assign a unique key for each master name server.
- TSIG key administration may be performed when modifying a secondary zone group.
- Assigning a TSIG key to a master name server will only authenticate communication generated for the zones associated with that secondary zone group. If other secondary zone groups also leverage the same master name server, then you will need to update their configuration as well.

- Changes made to a TSIG key are applied globally. All secondary zone groups that leverage that TSIG key will use the updated configuration when authenticating DNS communication.
- A TSIG key cannot be deleted if it is currently assigned to a master name server.

To add a TSIG key

1. Navigate to the **Route (DNS)** page.
2. Click on the secondary zone group that contains a master name server to which you would like to add or modify TSIG authentication.
3. Click **Manage TSIG Keys**.
4. In the **TSIG Alias** option, specify a brief, descriptive name for the TSIG key.
5. In the **Key Name** option, indicate a unique name by which the master and Route name servers will identify the key.
6. In the **Key Type** option, select the cryptographic hash function used to generate the key value.
7. In the **Key Value** option, paste the hash value generated for this key.
8. Click **Add**.
9. Click **Done**.
10. Find the desired master name server and assign it that TSIG key by selecting it from the **–Select TSIG –** option.
11. Click **Submit Group** to save your changes.

To modify a TSIG key

1. Navigate to the **Route (DNS)** page.
2. Click on the desired secondary zone group.
3. Click **Manage TSIG Keys**.
4. Click **Edit** next to the desired TSIG key.
5. Review the TSIG key and modify the alias, name, cryptographic hash function, and hash value as needed.
6. Click **Done** when finished.
7. Click **Done** to close TSIG management.
8. Click **Submit Group** to save your changes.

To delete a TSIG key

1. Navigate to the **Route (DNS)** page.
2. Click on the desired secondary zone group.
3. Click **Manage TSIG Keys**.
4. Click **Edit** next to the desired TSIG key.
5. Click the delete icon ( ).
6. Click **Done** to close TSIG management.
5. Click **Submit Group** to confirm the deletion of that TSIG key.

Secondary Zones

A secondary zone consists of:

- A read-only copy of a zone that resides on a master name server.
- Required read-only records that define vanity name servers.

Administration

Secondary zone management is performed through secondary zone groups. Each secondary zone group contains a list of zones. This list determines the set of zones that will be transferred to our DNS service. Once a zone has been transferred, it becomes a secondary zone within our DNS service.

A brief overview on administrative secondary zone tasks is provided below.

Task	Description
Add	A secondary zone can be created by simply adding the zone name to the desired secondary zone group. The zone will be transferred from the master name server upon saving the secondary zone group.
Modify	The zone name defined in the secondary zone group can be modified. This will cause our Route name servers to delete the old secondary zone and then perform a full zone transfer for the newly specified one.
Delete	A zone name can be deleted from a secondary zone group. Deleting a zone name will cause the deletion of the corresponding secondary zone.
View	A secondary zone can be viewed by clicking on it from the **Route (DNS)** page. Keep in mind that a secondary zone is read-only. All zone changes should be performed from the master name server. These changes will automatically be propagated to our network.

Secondary Zone Verification

Once you have added the desired secondary zones, it is important to verify that all of the records have been successfully transferred from the master name server to our Route name servers. This verification process consists of:

- Checking for transmission errors. This can be performed by clicking on the **Errors** link corresponding to the master name server that performed the zone transfer.
- Comparing the original zone to the secondary zone. You may view the raw data used to create the zone through the **Transfer Data** link in the **Master Server Group** section of your zone. Compare the raw data against the records imported for that zone.

Note: Each secondary zone indicates the serial number, master name server, TSIG, and the date/time stamp of when it was transferred to our system. Use this information when performing secondary zone verification.

DNSSEC

Secondary DNS supports zones that contain DNSSEC records. This allows DNSSEC records to be imported into the corresponding secondary zone.

DNSSEC Records

A brief description for the supported set of DNSSEC records is provided below.

Name	Description
RRSIG	Provides the DNSSEC signature through which DNS data is authenticated.
DNSKEY	Provides the public key through which a DNS resolver verifies the DNSSEC signature in a RRSIG record.
DS	Identifies a sub-delegated zone by its name. It also identifies a DNSKEY record in the sub-delegated zone.
NSEC	Indicates the next secured record in the zone by name. It also indicates the type of records in the zone that have been assigned that name. **Note:** A DNS resolver uses this record to verify that a record of a specific name and type does not exist within a zone.
NSEC3	Indicates the next secured record in the zone by hashed name. It also indicates the type of records in the zone that have been assigned that name. **Note:** A DNS resolver uses this record to verify that a record of a specific name and type does not exist within a zone.
NSEC3PARAM	Allows Authoritative DNS servers to determine the set of NSEC3 records to include in response to DNSSEC requests for a record that does not exist.

Name	Description
DLV	The DNSSEC Lookaside Validation registry record publishes DNSSEC trust anchors outside of the standard DNS delegation chain. This allows a DNS resolver to validate DNSSEC records via an alternative chain of trust.

Zone Transfer Authorization

Our Secondary DNS feature allows a zone to be transferred from one or more master name servers to our Route name servers. This initial zone transfer, along with subsequent synchronizations, requires that our Route transfer system be allowed access to the master name servers that host the primary zone.

Our Route transfer system communicates with master name servers over port 53 (UDP/TCP) using the following IP blocks:

IP blocks (IPv4):

```
152.195.162.0/27
192.16.60.0/24
```

IP blocks (IPv6):

```
2606:2800:4002:60::0/64
2606:2800:420f:1072::0/64
```

Enabling Secondary DNS

Creating a secondary zone group will not automatically redirect DNS traffic to our Route name servers. Once a secondary zone is ready to handle production traffic, its traffic needs to be directed to our Route name servers by registering our name servers at your domain registrar. The naming convention for our vanity name servers is:

s#ns#.<Zone Name>.

In the above naming convention, the pound symbol (#) represents a sequential number and <Zone Name> identifies the name of your zone.

Secondary Zone Delegation Example

This example indicates the records that will be created for a secondary zone called "secondexample.com."

Record Type	Name	Value
NS	@	s1ns1.secondexample.com.
NS	@	s2ns2.secondexample.com.
NS	@	s3ns3.secondexample.com.
NS	@	s4ns4.secondexample.com.
A	s1ns1	192.16.16.5
A	s2ns2	192.16.16.6
A	s3ns3	198.7.29.5
A	s4ns4	198.7.29.6
AAAA	s1ns1	2606:2800:3::5
AAAA	s2ns2	2606:2800:3::6
AAAA	s3ns3	2606:2800:c::5
AAAA	s4ns4	2606:2800:c::6

Authorize our Route service to perform secondary DNS tasks by registering all of our secondary zone name servers with your domain registrar.

Note: Secondary DNS support setup varies by domain registrar. Additional information can be found in the following Route Help Center article: FAQ: Route.