

# Scale Quickly and Efficiently with Dell ECS and Vertica Analytical Database

## Scalability for Data Warehouse Modernization and Data Lake Exploration

February 2023

H18684.1

## Reference Architecture

### Abstract

Vertica in Eon Mode separates compute from storage, allowing the data lake to scale seamlessly. Dell ECS provides exabyte scalability and extreme performance for object storage. By combining these technologies, Dell and Vertica offer customers a highly scalable, unified analytic platform for massive data lake exploration.

Dell Technologies Solutions

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA February 2023 H18684.1.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Contents

- Executive summary ..... 4**
- Introduction ..... 5**
- Key components ..... 6**
- Test environment ..... 7**
- Configuration and data loading ..... 9**
- Installation of Vertica in Eon Mode..... 10**
- Vertica test methodologies ..... 12**
- Best practices..... 13**
- References ..... 14**
- Appendix—GSLB configuration example ..... 15**

## Executive summary

### Overview

Vertica in Eon Mode gives companies operational flexibility for their analytic workloads by separating compute and storage. By leveraging S3 object storage, Vertica in Eon Mode allows organizations to scale resources with more flexibility than ever. Vertica in Eon Mode is available for on-premises deployments with Dell ECS, so analytics professionals and IT administrators can use cloud technology such as object stores while keeping data on the premises.

Vertica in Eon Mode with ECS gives customers the freedom to leverage cloud innovation for analytics wherever their data resides, without assuming the risks, costs, and complexities of cloud migration. By using Vertica in Eon Mode as your analytics engine and ECS as your object store, you can:

- Scale infrastructure resources independently. Storage can grow without the addition of expensive compute, and compute can be scaled up or down with variable or intermittent workloads.
- Isolate workloads. Business analysts and data scientists can work independently from a single source of truth without competing for resources.
- Simplify database options. Customers experience improved node recovery, superior workload balancing, and more while maintaining complete ACID compliance.
- Hibernate compute nodes. Customers can start and stop analytics more efficiently by hibernating compute nodes when the nodes are not needed.

The certification conducted for this solution is provided by Vertica and consists of several tests, including S3 API compatibility, backup and restore, and performance.

### Revisions

Date	Description
February 2021	Initial release
February 2023	Added EXF900 certification

### We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

**Authors:** Damien Mas, Rich Paulson

**Contributors:** Vertica Partner Engineering

---

**Note:** For links to additional ECS documentation, see the [ECS Info Hub](#).

---

## Introduction

### Introduction to Vertica

Vertica is a unified analytics data warehouse, based on a massively scalable architecture with the broadest set of analytical functions that span event and time series, pattern matching, geospatial, and end-to-end in-database machine learning. Vertica easily applies these powerful functions to the largest and most demanding analytical workloads, arming customers with predictive business insights.

#### Two modes, one unified platform

Vertica offers two deployment options, based on the same core platform, offering a choice of infrastructure. Vertica in Enterprise Mode runs on industry-standard servers with tightly coupled storage, delivering the highest performance for use cases that demand consistent compute capacity. Vertica in Eon Mode has a cloud-native architecture that separates compute from storage, enabling simplified management for variable workloads with the flexibility to apply specific compute resources to shared storage for different business use cases. By decoupling compute and storage, Vertica in Eon Mode allows compute and storage resources to scale independently without constraints, enabling a future-proof architecture.

### Introduction to ECS

Dell ECS is a software-defined, cloud-scale, object storage platform. With ECS, any organization can deliver scalable public cloud services with the reliability and control of a private-cloud infrastructure.

ECS provides comprehensive protocol support for unstructured—object and file—workloads on a single modern storage platform. Using ECS, organizations can easily manage globally distributed storage infrastructure under a single global namespace with anywhere access to content.

ECS features a flexible software-defined architecture that is layered to promote limitless scalability. Each layer is abstracted and independently scalable with high availability and no single points of failure. ECS also comes in a fully integrated turnkey appliance that bundles software and Dell PowerEdge servers into an easily deployed object system. The EX-Series, which is the third generation of ECS hardware appliances, builds on the legacy of two object storage platforms that predate ECS—Dell Centera and Dell Atmos. The ECS EX-Series consists of four unique hardware products: EX300, EX500, EX3000, and EXF900, which is an all NVMe solution.

### Terminology

The following terminology is used throughout this document:

**Table 1. Terminology**

Term	Description
Communal storage	<p>The shared storage location containing an Eon Mode database's data. The data within the communal storage is the canonical copy of the data—Vertica does not consider data as being committed until it has been written to communal storage.</p> <p>The communal data storage location is based on an object store, such as an S3 bucket in Dell ECS.</p>

Depot	A cache of data maintained by the nodes in an Eon Mode database to limit reads from communal storage. Retrieving data from communal storage often has high latency and potentially limited bandwidth, especially in cloud environments. Each node in the database caches the data it reads from and writes to the communal storage. When processing a query, the nodes first check this local cache for the data they need. If the data is cached locally, the nodes use the cached version instead of retrieving the data from communal storage.
-------	---

## Key components

### Vertica technology features

The following table provides a list of Vertica technology features used in the validation process.

**Table 2. Vertica technology features**

Technology model	Technology version	Description
Vertica in Eon Mode	10.0, 10.1.1-0	6-node physical servers
S3 API integration	Not applicable	Vertica in Eon Mode with S3 support

### Dell technology features

The following table provides a list of Dell technology features used in the validation process.

**Table 3. Dell ECS technology features**

Technology model	Technology version	Description
Dell EX500	3.6	Physical ECS 8-node appliance
Dell EXF900	3.7.0.2	Physical ECS 8-node appliance
S3 API protocol access	Not applicable	Support for a rich S3-compatible API
Dell PowerEdge R640 servers	CentOS 7	6x Dell PowerEdge R640, each with 2x Intel Xeon Gold 6254 CPU @ 3.1 GHz (2x 18 core/36 with HT each = 72 logical cores total with HT), 384 GB RAM, and Bonded 2x 25 GbE NICs  The Vertica Depot resides on 3x 1.6 NVMe (4.4 TB, RAID 0 - /dev/md0).
<ul style="list-style-type: none"> <li>Dell PowerSwitch S5148</li> <li>Dell PowerSwitch S5248</li> </ul>	Not applicable	An innovative, future-ready top-of-rack (ToR) open networking switch, providing excellent capabilities and cost-effectiveness

## Data access methods

The following table provides a list of data access methods and technologies used in the validation process.

**Table 4. Data access methods**

Data access method	Data access structure	Details
Integrated S3 API	API	Used the native S3 API present in both Dell and Vertica technologies

## Test environment

### Architecture overview

Figure 1 illustrates the architecture, showing the key components of the solution as it was tested and benchmarked.

---

**Note:** In a customer deployment, an IP load balancer to direct traffic among the ECS nodes is highly recommended. Dell recommends the use of Kemp ECS Connection Manager appliances for this purpose. The IP Load Balancer must support a throughput equivalent to the aggregated throughput of every Vertica Client running on the Dell PowerEdge servers.

During our tests, we used a DNS with round robin (DNS RR) instead of a load balancer. The DNS was responsible for redirecting the traffic across the ECS nodes. The use of DNS RR is not recommended in a production environment; if an ECS node goes down, the DNS will not be aware of the failure.

---

All testing appliances were contained within a single lab environment, except for the corporate DNS service.

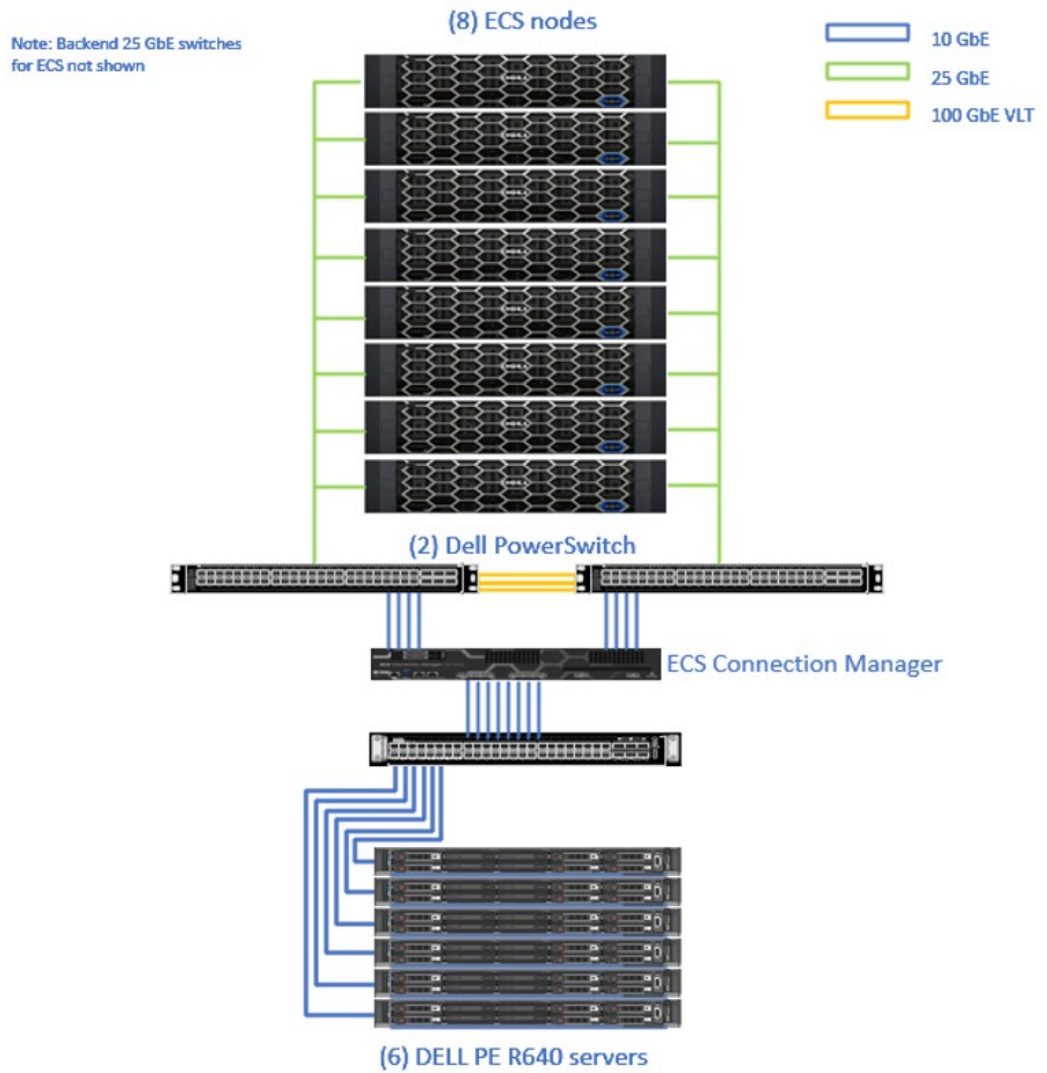


Figure 1. Dell ECS architecture overview

Vertica test software

The following table provides information about the Vertica test software.

Table 5. Vertica test software

Description	Detail
Vertica S3 API compatibility suite	Designed to work with AWS S3 compatible object store.
Vertica S3 benchmark suite	Based on TPC Benchmark DS (TPC-DS), which is a decision support benchmark that models several applicable aspects of a decision support system, including queries and data maintenance. A benchmark result measures query response time in single user mode; query throughput in multi-user mode; and data maintenance performance for a given hardware, operating system, and data processing system configuration under a controlled, complex, and multi-user decision support workload.



# Configuration and data loading

## Introduction

ECS supports both traditional object users and IAM users. The examples in this section use a traditional object user. An object user (access key ID), S3 secret key, namespace, and bucket must be created for Vertica. A VIP is also required on an IP load balancer to direct traffic across the ECS nodes. Dell recommends using Kemp ECS Connection Manager. For more details about the configuration of the ECS Connection Manager load balancer, see [Dell ECS with Kemp ECS Connection Manager](#).

## Create namespace

Create a namespace as follows:

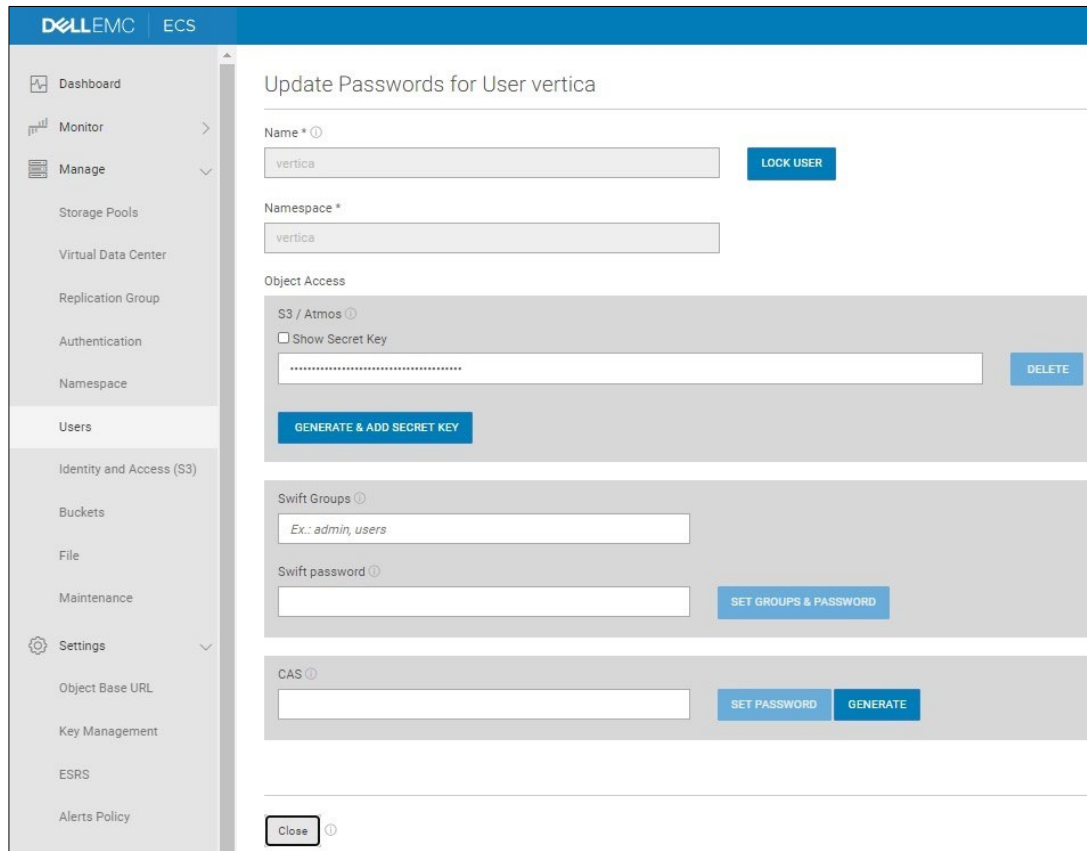
1. Log in to the ECS web management console.
2. Select **Manage > Namespace**.
3. Enter a namespace in the **Name** field as shown, and then click **Save**.

## Create object user

Create an object user as follows:

1. Select **Manage > Users**.
2. Enter a username in the **Name** field.
3. Click **NEXT TO ADD PASSWORD**.

4. Click **GENERATE & ADD SECRET KEY** to generate a secret key for the new object user.



5. Select **Show Secret Key** to retrieve the secret. You will need this secret later when configuring Vertica.

## Installation of Vertica in Eon Mode

### Introduction

For details about installing and configuring Vertica for in Eon Mode with the s3 object storage, see the [Vertica documentation](#).

---

**Note:** This document assumes that Vertica has been installed on the client systems.

---

## Prepare the environment

Prepare the environment as follows:

1. On the first Vertica client, log in as `dbadmin` user, which was created during Vertica installation.
2. Set up the ECS credentials with the AWS CLI tool by running the following command:

```
aws configure
AWS Access Key ID [None]: vertica
AWS Secret Access Key [None]: <yoursecretkey>
Default region name [None]: us-east-1
Default output format [None]:
```

3. Create the `auth_params.conf` file in the `dbadmin` user's home directory as follows:

```
[dbadmin@vertica-client-1 ~]$ cat auth_params.conf
awsauth = <youraccesskey>:<yoursecretkey>
awsendpoint = s3.<ecsendpointexample>.com:9020
awsregion = us-east-1
awsenablehttps = 0
```

4. Create the necessary buckets on ECS by running the following command:

```
aws s3 mb --endpoint-url=http://s3.<ecsendpointexample>.com:9020 s3://objstorc
aws s3 mb --endpoint-url=http://s3.<ecsendpointexample>.com:9020 s3://objstorb
aws s3 mb --endpoint-url=http://s3.<ecsendpointexample>.com:9020 s3://objstore
```

5. List the buckets to verify that they were created:

```
aws s3 ls --endpoint-
url=http://s3.<ecsendpointexample>.com:9020
```

6. Verify the expected output:

```
2020-10-13 10:27:18 objstorb
2020-10-13 10:26:49 objstorc
2020-10-13 10:28:00 objstore
```

7. Run the following S3 API test command:

```
./s3api.sh | tee ecs_test.txt
```

## Create the Vertica in Eon Mode database

Follow these steps to create the Vertica test database and ensure functionality:

1. Create the database:

```
admintools -t create_db -x auth_params.conf -s <add database server IPs example: x.x.x.x,x.x.x.x,x.x.x.x,x.x.x.x,x.x.x.x >
-d VMart -p databasepwd --depot-path=/data/db/mydepot --shard-count=12 --communal-storage-location=s3://objstorc -D /home/dbadmin -c /home/dbadmin -l /home/dbadmin/DELL.dat --depot-size 500GB
```

2. Run vsql to change the dbadmin's password to empty string:

```
vsql -U dbadmin -w databasepwd
> alter user dbadmin identified by '';
```

3. Alter the database default setting for AWSStreamingConnectionPercentage:

```
> ALTER DATABASE DEFAULT SET PARAMETER
AWSStreamingConnectionPercentage = 0;
```

4. Activate native load balancing:

```
> SELECT SET_LOAD_BALANCE_POLICY('ROUNDROBIN');
```

5. Create a small table, insert a row, and query the table to ensure basic functionality:

```
> create table test(c1 int);
> insert into test values('1');
> commit;
> select * from test;
> drop table test cascade;
# exit vsql
\q
```

## Vertica test methodologies

### Introduction

The Vertica test methodologies followed Vertica's certification requirements for third-party partners and vendors. The test methodologies included compatibility testing, data protection and recovery testing, and performance testing.

### ECS compatibility testing

The following table describes the test to validate Vertica in Eon Mode platform functionality when the platform is integrated with Dell ECS as the on-premises cloud storage appliance.

**Table 6. S3 API compatibility test**

Test	Description
S3 API compatibility	This test performs several API calls to verify compatibility with the object store under test. It includes requests such as PUT, GET, LIST objects, and multi-part uploads.

## ECS data protection and recovery testing

The Vertica test suite includes a series of tests that focus on database protection and recovery. These tests include database backup, database restore, and database revive tests, as described in the following table.

**Table 7. Data protection and recovery tests**

Test	Description
Database backup test	The backup test makes a backup of the database to a different bucket, like a snapshot in time.
Database restore test	The restore test restores a database that was previously backed up.
Database revive test	The revive test simulates that you lost your physical hosts and built a replacement cluster. In EON mode with separation for compute and storage, you can change the hardware and then revive the database from communal storage to the new hardware.

## ECS performance testing

The Vertica test suite also includes a series of TPC-DS benchmarks that measure queries that mimic a decision support system.

**Table 8. Performance tests**

Test	Description
Single-user query tests	These tests are run using a 1 TB and 10 TB dataset with the Depot mode set to on as well off.
Four-user query tests	These tests are run using a 1 TB and 10 TB dataset with the Depot mode set to on as well off.

## ECS test results

All the tests ran successfully against both ECS models and passed Vertica's stringent requirements.

## Best practices

Consider the following recommendations when using Vertica in Eon Mode with Dell ECS:

- **IP load balancer**—Configure the load balancer in GSLB mode to increase throughput. With ECS Connection Manager, GSLB has the added benefit of performing health checks on the ECS nodes.  
For an example of ECS Connection Manager GSLB configuration, see [Appendix—GSLB configuration example](#).
- **SSD cache disk**—Add an SSD cache disk for metadata caching. The addition of this disk will improve system-wide read latency and transactions per second (TPS) for small files.
- **Depot mode**—In most cases, configure Vertica EON in Depot On mode.

## References

### Dell Technologies documentation

The following links provide additional information about ECS. Access to documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [Dell Technologies ECS Info Hub](#)
- [Dell Technologies Storage Hub](#)

### Vertica documentation

Vertica documentation can be accessed publicly:

- [Vertica documentation](#)

## Appendix—GSLB configuration example

### Introduction

In environments where a single appliance cannot deliver the necessary throughput or transactions per second, clustering can be used to scale the load balancers. The most effective way to accomplish this level of scale is to deploy the load balancers in a Global Server Load Balancer (GSLB) cluster (sometimes referred to as a GEO cluster). This approach scales linearly when more load balancers are added to the cluster, with combined resources delivering the required performance.

GSLB clustering uses DNS to distribute traffic across the load balancers. As clients and other systems attempt a connection to the published application, DNS delegation is used to direct the resolution of the Fully Qualified Domain Name (FQDN) to GSLB on the load balancers. GSLB returns an IP address to the client based on the scheduling method defined in the configuration, providing proper and efficient distribution across the cluster, thus using each node in each ECS VDC.

### Create the GSLB cluster

A cluster entry must be created for each load balancer that will participate in the cluster. These entries are used within the FQDN configuration to enable the advanced health-checking capabilities.

---

**Note:** This process only needs to be performed on one load balancer in the cluster. Settings are propagated to other members.

---

Go to **Global Balancing > Manage Clusters** and add the IP address and name for each load balancer in the cluster.

Configured Clusters					
IP Address	Name	Coordinates	Type	Checker	Availability
10.246.176.29	ecsmgrh3-01	0°0'0"N 0°0'0"W	Remote LM	Implicit	Up
10.246.176.30	ecsmgrh3-02	0°0'0"N 0°0'0"W	Remote LM	Implicit	Up

### Enable GEO synchronization

To facilitate management of the GSLB configuration, enable GEO synchronization on the load balancers in the cluster. Entering the settings on a single load synchronizes them across all other members in the cluster.

To configure GEO synchronization, go to **Certificates & Security > Remote Access**.

GEO Settings	
Remote GEO ECS Connection Manager Access	10.246.176.29 10.246.176.30 <a href="#">Set GEO ECS Connection Manager access</a>
GEO ECS Connection Manager Partners	10.246.176.30 <a href="#">Set GEO ECS Connection Manager Partners</a>
Partner Status	10.246.176.30
GEO ECS Connection Manager Port	22 <a href="#">Set GEO ECS Connection Manager Port</a>
GEO Update Interface	eth0: 10.246.176.29

## Create the GSLB FQDN

An FQDN is the URL that clients and other systems will use to access ECS. The distribution method is defined and health-checking is specified within the GSLB FQDN configuration.

Create the GSLB FQDN as follows.

**Note:** The following steps only need to be performed on one load balancer in the cluster. Settings are propagated to other members.

1. Go to **Global Balancing > Manage FQDNs**, add the FQDN, and then click **Add FQDN**.

Configured Fully Qualified Domain Names

New FQDN  **Add FQDN**

2. Set the distribution method to **Round Robin**.

<-Back

Configure s3.example.com

Selection Criteria

Public Requests

Private Requests

Site Failure Handling Failure Delay (minutes)  **Set Failure Delay**

Enable Local Settings

Unanimous Cluster Health Checks

3. Expand the **IP Address** section and enter the IP address of an ECS node. Do not select a cluster from the **Cluster** menu.

IPAddresses

New IP Address  Cluster  **Add Address**

IP Address	Cluster	Checker	Availability	Parameters	Operation
------------	---------	---------	--------------	------------	-----------

4. Click **Add Address**, and then continue adding the IP addresses for the other nodes in the ECS virtual data center.
5. After adding all the ECS nodes, configure the health check for each entry.

**Note:** This step can also be done immediately after you add the IP address for an ECS node.

- a. Modify the **Checker** dropdown to **HTTP**.
- b. At **Set Address**, add the IP address of the ECS node, change the port to 9020, and then click **Set Address**.
- c. **At Set URL**, enter `/?ping`, and then click **Set URL**.



- d. Leave the **Set Status Codes** and **Set Host** fields empty.
- e. Leave the HTTP request method as **GET**.

IP Address	Cluster	Checker	Availability	Parameters	Operation
10.246.22.171	Select Cluster ▾	HTTP ▾			Disable   Delete
		10.246.22.171   9020   Set Address			
		/?ping   Set URL	✓ Up		
		Set Status Codes			
		Set Host			
		GET ▾			

## Set GSLB parameters

Several parameters must be configured for GSLB clustering to operate efficiently. These settings reflect general DNS best practices and ensure proper distribution of traffic.

1. Go to **Global Balancing > Miscellaneous Params**.
2. Under **Zone**, enter the **Zone Name** for the environment.
3. Under **Source of Authority**, enter the responses to be sent for each of these requests:
  - **Source of Authority** (response sent for SOA requests)
  - **Name Server** (response sent for NS requests)
  - **SOA Email** (email response sent for SOA requests)
4. Under **Global**, enter **1** for **TTL**.
 

This is the time to live setting, which reduces the DNS resolver cache to ensure even distribution across the load balancers in the cluster. This can also be set within the FQDN settings.
5. Under **Resource Check Parameters**, enter the following information to ensure that the health of the systems is reported near real time:
  - a. For **Check Interval**, enter **9** and click **Set**.
  - b. For **Connection Timeout**, enter **4** and click **Set**.
  - c. For **Retry attempts**, enter **2** and click **Set**.
6. Ensure **Stickiness** is set to 0.
 

This setting ensures that GSLB will not send the same response back to the client for the specified amount of time.

## Set up DNS delegation

Delegate the zone to the ECS Connection Manager GSLB. The DNS environment essentially delegates the FQDN to each load balancer in the GSLB cluster.

---

**Note:** The steps to configure the delegation vary, based on different DNS environments.

---