



# Pulse Secure Virtual Traffic Manager: Terraform Provider Reference Guide

Supporting Pulse Secure Virtual Traffic Manager 20.3

Product Release **20.3**

Published **8 February, 2021**

Document Version **1.0**

Pulse Secure, LLC  
2700 Zanker Road,  
Suite 200 San Jose  
CA 95134

[www.pulsesecure.net](http://www.pulsesecure.net)

© 2021 by Pulse Secure, LLC. All rights reserved.

Pulse Secure and the Pulse Secure logo are trademarks of Pulse Secure, LLC in the United States. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Pulse Secure, LLC assumes no responsibility for any inaccuracies in this document. Pulse Secure, LLC reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Pulse Secure Virtual Traffic Manager: Terraform Provider Reference Guide*

The information in this document is current as of the date on the title page.

## **END USER LICENSE AGREEMENT**

The Pulse Secure product that is the subject of this technical documentation consists of (or is intended for use with) Pulse Secure software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at <http://www.pulsesecure.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

# Contents

---

PREFACE .....	1
DOCUMENT CONVENTIONS .....	1
TEXT FORMATTING CONVENTIONS.....	1
COMMAND SYNTAX CONVENTIONS.....	1
NOTES AND WARNINGS.....	2
REQUESTING TECHNICAL SUPPORT .....	2
SELF-HELP ONLINE TOOLS AND RESOURCES.....	2
OPENING A CASE WITH PSGSC .....	3
OVERVIEW .....	5
ABOUT THIS GUIDE.....	5
INTRODUCING THE TRAFFIC MANAGER.....	5
INTRODUCING THE TERRAFORM PROVIDER .....	5
GETTING STARTED.....	7
INSTALLING THE TRAFFIC MANAGER TERRAFORM PROVIDER.....	7
INITIALIZING THE TERRAFORM PROVIDER .....	7
RESOURCE AND DATA SOURCE REFERENCE .....	9
ABOUT THE RESOURCE MODEL.....	9
USING THE RESOURCE MODEL.....	9
CONFIGURATION RESOURCES .....	10
SETTING CONFIGURATION VALUES USING THE RESOURCE MODEL .....	10
DEPENDANCY TREE.....	11
RESOURCE REFERENCE .....	11
VTM_ACTION_PROGRAM.....	11
VTM_SSL_ADMIN_CA .....	12
VTM_ACTION .....	12
VTM_APTIMIZER_SCOPE.....	16
VTM_BGPNEIGHBOR.....	17
VTM_BANDWIDTH.....	18
VTM_CLOUD_API_CREDENTIAL.....	19
VTM_CUSTOM .....	20
VTM_DNS_SERVER_ZONE .....	21
VTM_DNS_SERVER_ZONE_FILE.....	21
VTM_EVENT_TYPE.....	22
VTM_EXTRA_FILE.....	25

VTM_GLB_SERVICE	26
VTM_GLOBAL_SETTINGS	30
VTM_KERBEROS_KRB5CONF	62
VTM_KERBEROS_KEYTAB	63
VTM_KERBEROS_PRINCIPAL	63
VTM_LICENSE_KEY	64
VTM_LOCATION	65
VTM_LOG_EXPORT	66
VTM_MONITOR	67
VTM_MONITOR_SCRIPT	72
VTM_APPLIANCE_NAT	72
VTM_POOL	74
VTM_PROTECTION	89
VTM_APPLICATION_FIREWALL	92
VTM_RATE	92
VTM_RULE	93
VTM_SERVICE_LEVEL_MONITOR	93
VTM_SSL_CLIENT_KEY	94
VTM_SSL_SERVER_KEY	95
VTM_SSL_TICKET_KEY	96
VTM_SSL_CA	97
VTM_SECURITY	98
VTM_SERVICEDISCOVERY	99
VTM_PERSISTENCE	99
VTM_TRAFFIC_IP_GROUP	101
VTM_TRAFFIC_MANAGER	104
VTM_RULE_AUTHENTICATOR	118
VTM_SAML_TRUSTEDIDP	119
VTM_USER_AUTHENTICATOR	120
VTM_USER_GROUP	125
VTM_VIRTUAL_SERVER	126
VTM_OPTIMIZER_PROFILE	155
DATA SOURCES	156
RESOURCES	156
RESOURCE LISTS	157
STATISTICS	158
VTM_ACTION_STATS	159
VTM_CACHE_ASP_SESSION_CACHE_STATS	159
VTM_BANDWIDTH_STATS	160
VTM_CLOUD_API_CREDENTIAL_STATS	160
VTM_CONNECTION_RATE_LIMIT_STATS	161

VTM_EVENT_STATS.....	162
VTM_GLB_SERVICE_STATS.....	162
VTM_GLOBALS_STATS.....	163
VTM_TRAFFIC_IPS_IP_GATEWAY_STATS.....	170
VTM_CACHE_IP_SESSION_CACHE_STATS.....	171
VTM_CACHE_J2EE_SESSION_CACHE_STATS.....	172
VTM_LISTEN_IP_STATS.....	173
VTM_LOCATION_STATS.....	173
VTM_NETWORK_INTERFACE_STATS.....	174
VTM_NODES_NODE_STATS.....	175
VTM_NODES_NODE_INET46_STATS.....	177
VTM_PER_LOCATION_SERVICE_STATS.....	178
VTM_PER_NODE_SLM_PER_NODE_SERVICE_LEVEL_STATS.....	179
VTM_PER_NODE_SLM_PER_NODE_SERVICE_LEVEL_INET46_STATS.....	180
VTM_NODES_PER_POOL_NODE_STATS.....	181
VTM_POOL_STATS.....	183
VTM_RULE_AUTHENTICATOR_STATS.....	186
VTM_RULE_STATS.....	187
VTM_SERVICE_LEVEL_MONITOR_STATS.....	187
VTM_SERVICE_PROTECTION_STATS.....	188
VTM_CACHE_SSL_CACHE_STATS.....	190
VTM_SSL_OCSP_STAPLING_STATS.....	191
VTM_CACHE_SSL_SESSION_CACHE_STATS.....	192
VTM_TRAFFIC_IPS_TRAFFIC_IP_STATS.....	192
VTM_TRAFFIC_IPS_TRAFFIC_IP_INET46_STATS.....	193
VTM_CACHE_UNI_SESSION_CACHE_STATS.....	194
VTM_EXTRAS_USER_COUNTERS_32_STATS.....	195
VTM_EXTRAS_USER_COUNTERS_64_STATS.....	195
VTM_VIRTUAL_SERVER_STATS.....	196
VTM_CACHE_WEB_CACHE_STATS.....	203
SYSTEM INFORMATION.....	205
VTM_BACKUPS_FULL.....	205
VTM_INFORMATION.....	206
VTM_STATE.....	206
TABLE FIELD DATA SOURCES.....	208
EXAMPLE DEPLOYMENT TEMPLATE.....	211



# Preface

---

- [Document conventions](#) ..... 1
- [Requesting Technical Support](#) ..... 2

## Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Pulse Secure technical documentation.

### Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

Format	Description
<b>bold text</b>	Identifies command names
	Identifies keywords and operands
	Identifies the names of user-manipulated GUI elements
	Identifies text to enter at the GUI
<i>italic text</i>	Identifies emphasis
	Identifies variables
	Identifies document titles
Courier Font	Identifies command output
	Identifies command syntax examples

### Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
<b>bold text</b>	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.

Convention	Description
{ <b>x</b>   <b>y</b>   <b>z</b> }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
<b>x</b>   <b>y</b>	A vertical bar separates mutually exclusive elements.
< >	Non-printing characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, member[member...].
\	Indicates a “soft” line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

## Notes and Warnings

Note, Attention, and Caution statements might be used in this document.

**Note:** A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

### ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

### CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

## Requesting Technical Support

Technical product support is available through the Pulse Secure Global Support Center (PSGSC). If you have a support contract, file a ticket with PSGSC.

- Product warranties—For product warranty information, visit <https://support.pulsesecure.net/product-service-policies/>

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Pulse Secure provides an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://support.pulsesecure.net>
- Search for known bugs: <https://support.pulsesecure.net>
- Find product documentation: <https://www.pulsesecure.net/techpubs>
- Download the latest versions of software and review release notes: <https://support.pulsesecure.net>
- Open a case online in the CSC Case Management tool: <https://support.pulsesecure.net>



- To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://support.pulsesecure.net>

For important product notices, technical articles, and to ask advice:

- Search the Pulse Secure Knowledge Center for technical bulletins and security advisories: <https://kb.pulsesecure.net>
- Ask questions and find solutions at the Pulse Community online forum: <https://community.pulsesecure.net>

## Opening a Case with PSGSC

You can open a case with PSGSC on the Web or by telephone.

- Use the Case Management tool in the PSGSC at <https://support.pulsesecure.net>.
- Call 1-844 751 7629 (Toll Free, US).

For international or direct-dial options in countries without toll-free numbers, see <https://support.pulsesecure.net/support/support-contacts/>



# Overview

---

This chapter provides an overview of Pulse Secure Virtual Traffic Manager (the Traffic Manager). This chapter contains the following sections:

- [About This Guide](#) ..... 5
- [Introducing the Traffic Manager](#) ..... 5
- [Introducing the Terraform Provider](#) ..... 5

## About This Guide

The *Pulse Secure Virtual Traffic Manager: Terraform Provider Reference Guide* describes how to use the Terraform provider to create and configure Traffic Manager instance through deployment templates.

Read this guide for an introduction to the functionality available as part of the provider, for a complete reference to all Traffic Manager configuration resources and data sources, and for examples of how to use Terraform to provision your Traffic Manager deployment.

For a detailed description of the Traffic Manager and its full feature set, see the *Pulse Secure Virtual Traffic Manager: User's Guide*.

## Introducing the Traffic Manager

The Traffic Manager product family provides high-availability, application-centric traffic management and load balancing solutions in a range of software, hardware-ready, virtual appliance, and cloud-compute product variants. They provide control, intelligence, security and resilience for all your application traffic.

The Traffic Manager is intended for organizations hosting valuable business-critical services, such as TCP-based and UDP-based services like HTTP (web) and media delivery, and XML-based services such as Web Services.

## Introducing the Terraform Provider

Terraform is an open-source infrastructure-as-code tool produced by HashiCorp, Inc.

System administrators use Terraform to create and maintain configuration templates for your application infrastructure. Pulse Secure provides a set of tools to Terraform, known as a *provider*, that enable you to define Traffic Manager configurations within your application templates. Terraform can then provision Traffic Manager instances as part of a wider deployment.

To learn more about Terraform, see <https://www.terraform.io>.



# Getting Started

This chapter contains information about getting started using the Traffic Manager Terraform provider. This chapter contains the following sections:

- [Installing the Traffic Manager Terraform Provider](#) ..... 7
- [Initializing the Terraform Provider](#) ..... 7

## Installing the Traffic Manager Terraform Provider

To obtain the Pulse Secure Virtual Traffic Manager Terraform provider, download the source files from the *pulse-vadc* pages at GitHub:

<https://github.com/pulse-vadc/terraform-provider-vtm>

Make sure you download the version that corresponds to your target Traffic Manager's REST API version.

To install the provider in Terraform, follow the third-party plugin instructions at:

<https://www.terraform.io/docs/configuration/providers.html#third-party-plugins>

## Initializing the Terraform Provider

To use the Traffic Manager Terraform provider, initialize it within your Terraform template. The provider requires four arguments:

Field Name	Type	Description
base_url	string	The base URL of the Traffic Manager REST API. Use either a direct Traffic Manager URL, or through a Pulse Secure Services Director (Services Director) proxy.
username	string	The admin username of either the Traffic Manager or the Services Director proxy you are connecting to.
password	string	The admin password of either the Traffic Manager or the Services Director proxy you are connecting to.
verify_ssl_cert	boolean	If set to true, the provider checks that the Traffic Manager/Services Director provides a trusted SSL certificate when connecting to the REST API. If set to false, untrusted (for example, self-signed) certificates are accepted. Default: true.

The following code provides examples of how to initialize the provider:

```
# Direct to the Traffic Manager
provider "vtm" {
  base_url = "https://192.0.2.10:9070/api"
```

```
    username = "admin"
    password = "mYvtm53cret!"
    verify_ssl_cert = true
}

# Through a Services Director proxy
provider "vtm" {
    base_url = "https://192.0.2.100:8100/api/tmcm/<VERSION>/instance/<INSTANCE_ID>"
    username = "admin"
    password = "mYsd53cret!"
    verify_ssl_cert = true
}
```

# Resource and Data Source Reference

---

This chapter introduces the resource model, demonstrates its usage through the provider, and includes a complete listing of all available Traffic Manager configuration resources and data sources. This chapter contains the following sections:

- [About the Resource Model](#) ..... 9
- [Using the Resource Model](#) ..... 9
- [Configuration Resources](#) ..... 10
- [Data Sources](#) ..... 156

## About the Resource Model

The resource model reference included in this chapter lists all configuration resources and data sources available through the Traffic Manager REST API. Each sub-section relates to a specific resource type or data source, showing its name, description, unique URI path, and a table of properties.

Each property contains a description and data type. Additional information is provided where applicable, such as default value and permitted values (for enumerated types). For Table-type properties, a list of the Primary and Sub keys is provided.

For configuration resources, each property, and table keys where applicable, include an indicator of whether or not that item is a required value when creating configuration based on that resource.

For further information concerning the Traffic Manager REST API, see the *Pulse Secure Virtual Traffic Manager: REST API Guide*, available from the Pulse Secure website.

## Using the Resource Model

Resources can represent objects of a singular nature (for example, global settings) where only one instance of that resource type is ever present, or collections of objects (for example, virtual servers) where multiple instances might be created. To uniquely identify a specific instance of a collection-type resource, use the "name" property. Resources of a singular type do not require the "name" property.

Furthermore, there are two types of collection:

- Key/value resources. For example, virtual servers.
- Raw text files. For example, TrafficScript rules

Key/value resources have unique data structures, as defined in the resource model reference, whereas raw file resources have only two parameters; "name" and "content".

**ATTENTION**

Pulse Secure strongly recommends that singular object resources are only specified in one template for any cluster. Specifying these resources in multiple templates can lead to altering cluster-wide settings that impact other services.

## Configuration Resources

Configuration resources map to objects in the Traffic Manager's configuration system. The Terraform provider uses these resource definitions to provision services hosted on your Traffic Manager instances.

## Setting Configuration Values Using the Resource Model

The following code samples show how to set configuration properties for different resource types:

```
# Single object configuration
resource "vtm_global_settings" "global_settings" {
  admin_support_tls1 = false
  connection_idle_timeout = 30
}

# Collection-type key/value object configuration
resource "vtm_pool" "my_pool" {
  name = "MyPool"
  monitors = ["Ping", "Simple HTTP"]
  load_balancing_algorithm = "weighted_least_connections"
  nodes_table {
    node = "192.0.2.101:80"
    weight = 1
  }
  nodes_table {
    node = "192.0.2.102:80"
    weight = 3
  }
  nodes_table {
    node = "192.0.2.103:80"
    weight = 3
    state = "disabled"
  }
}

# Collection-type raw text object configuration
resource "vtm_rule" "my_rule" {
  name = "MyRule"
  content = <<EOF
$hostHeader = http.getHostHeader();

if($hostHeader == "www.example.com") {
  pool.use("ExamplePool");
}
EOF
```



```
}
```

## Dependency Tree

Generally speaking, the Traffic Manager REST API is tolerant of objects being created out-of-order. In other words, the REST API allows you to create a virtual server that references a TrafficScript rule that does not yet exist, even though this would trigger an error in the Traffic Manager diagnostics and event log. This error is suppressed after the corresponding rule is subsequently added. However, there are exceptions to this arrangement.

Pulse Secure recommends that all resource dependencies are explicitly specified by using parameter interpolation when creating relationships, rather than hard-coding names. For example, rather than providing a list of rule names to a virtual server as strings, specify the names as variables to aid Terraform in building the correct dependency tree. The following code samples demonstrate this process.

The code sample below shows the partial configuration of a virtual server with named request rules:

```
resource "vtm_virtual_server" "example" {
  name = "ExampleVS"
  request_rules = ["rule1", "rule2"]
  ...
}
```

Rather than specifying the hard-coded strings "rule1" and "rule2" in the list of request rules, use instead a set of variables to provide the rule names. The following code sample demonstrates the syntax used to reference properties of other resource objects in the provider:

```
resource "vtm_virtual_server" "example" {
  name = "ExampleVS"
  request_rules = ["${vtm_rule.rule1.name}", "${vtm_rule.rule2.name}"]
  ...
}
```

## Resource Reference

Use the configuration resource definitions in this section to provision services on your Traffic Manager instances.

### **vtm\_action\_program**

URI Endpoint: /api/tm/8.1/config/active/action\_programs

This is a program or script that can be referenced and used by actions of type 'Program'

```
# Example usage of vtm_action_program resource...
resource "vtm_action_program" "my_action_program" {
  content = ""
  name = "MyActionProgram"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_ssl\_admin\_ca

URI Endpoint: /api/tm/8.1/config/active/ssl/admin\_cas

The conf/ssl/admin\_cas directory contains SSL certificate authority certificates (CAs) and certificate revocation lists (CRLs) which can be used when validating connections made by the admin server for user authentication. CAs and CRLs can be managed under the Catalogs > SSL > Admin CAs and CRLs section of the Admin Server UI or by using functions under the Catalog.SSL.AdminCertificateAuthorities section of the SOAP API and CLI.

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_action

URI Endpoint: /api/tm/8.1/config/active/actions

A response to an event occurring in your traffic manager. An example of an action might be sending an email or writing a line to a log file.

```
# Example usage of vtm_action resource...
resource "vtm_action" "my_action" {
  name = "MyAction"
  type =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
email_from	The e-mail address from which messages will appear to originate. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "VTM@%hostname%"</li> </ul>
email_server	The SMTP server to which messages should be sent. This must be a valid IPv4 address or resolvable hostname (with optional port). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
email_to	A set of e-mail addresses to which messages will be sent. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
log_file	The full path of the file to log to. The text %zeushome% will be replaced with the location where the software is installed. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	A description of the action. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
program_arguments	A table containing arguments and argument values to be passed to the event handling program. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): The name of the argument to be passed to the event handling program. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>value (String): The value of the argument to be passed to the event handling program. (Required)</li> <li>description (String): A description for the argument provided to the program.</li> </ul> </li> </ul>

Property	Description
program_arguments_json	<p>A Traffic Manager REST-compatible JSON representation of the "program_arguments" table property. Use this field with the "vtm_action_arguments_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
program_program	<p>The program to run.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
soap_additional_data	<p>Additional information to send with the SOAP call.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
soap_password	<p>The password for HTTP basic authentication.</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
soap_proxy	<p>The address of the server implementing the SOAP interface (For example, https://example.com).</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
soap_username	<p>Username for HTTP basic authentication. Leave blank if you do not wish to use authentication.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
syslog_msg_len_limit	<p>Maximum length in bytes of a message sent to the remote syslog. Messages longer than this will be truncated before they are sent.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2048"</li> </ul>
syslog_sysloghost	<p>The host and optional port to send syslog messages to (if empty, messages will be sent to localhost).</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
timeout	<p>How long the action can run for before it is stopped automatically (set to 0 to disable timeouts).</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>

Property	Description
trap_auth_password	The authentication password for sending a Notify over SNMPv3. Blank to send unauthenticated traps. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trap_community	The community string to use when sending a Trap over SNMPv1 or a Notify over SNMPv2c. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trap_hash_algorithm	The hash algorithm for SNMPv3 authentication. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "md5"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"md5": MD5</li> <li>"sha1": SHA-1</li> </ul> </li> </ul>
trap_priv_password	The encryption password to encrypt a Notify message for SNMPv3. Requires that authentication also be configured. Blank to send unencrypted traps. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trap_traphost	The hostname or IPv4 address and optional port number that should receive traps. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trap_username	The SNMP username to use to send the Notify over SNMPv3. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trap_version	The SNMP version to use to send the Trap/Notify. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "snmpv1"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"snmpv1": SNMPv1</li> <li>"snmpv2c": SNMPv2c</li> <li>"snmpv3": SNMPv3</li> </ul> </li> </ul>

Property	Description
type	The action type. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"email": E-Mail</li> <li>"log": Log to File</li> <li>"program": Program</li> <li>"soap": SOAP Callback</li> <li>"syslog": Log to Syslog</li> <li>"trap": SNMP Notify or Trap</li> </ul> </li> </ul>
verbose	Enable or disable verbose logging for this action. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

### vtm\_optimizer\_scope

URI Endpoint: /api/tm/8.1/config/active/optimizer/scopes

Application scopes define criteria that match URLs to specific logical web applications hosted by a virtual server.

```
# Example usage of vtm_optimizer_scope resource...
resource "vtm_optimizer_scope" "my_optimizer_scope" {
  name = "MyOptimizerScope"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
canonical_hostname	If the hostnames for this scope are aliases of each other, the canonical hostname will be used for requests to the server. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
hostnames	The hostnames to limit acceleration to. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
root	The root path of the application defined by this application scope. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "/"</li> </ul>

### vtm\_bgpneighbor

URI Endpoint: /api/tm/8.1/config/active/bgpneighbors

The conf/bgpneighbors directory contains configuration files for BGP neighbors. The name of a file is the name of the neighbor configuration that it defines. BGP neighbors can be managed under the System > Fault Tolerance > BGP Neighbors section of the Admin UI, or by using functions under the BGPNeighbors section of the SOAP API and CLI.

```
# Example usage of vtm_bgpneighbor resource...
resource "vtm_bgpneighbor" "my_bgpneighbor" {
  name = "MyBgpneighbor"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
address	The IP address of the BGP neighbor <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
advertisement_interval	The minimum interval between the sending of BGP routing updates to neighbors. Note that as a result of jitter, as defined for BGP, the interval during which no advertisements are sent will be between 75% and 100% of this value. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
as_number	The AS number for the BGP neighbor <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65534"</li> </ul>

Property	Description
authentication_password	The password to be used for authentication of sessions with neighbors <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
holdtime	The period after which the BGP session with the neighbor is deemed to have become idle - and requires re-establishment - if the neighbor falls silent. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "90"</li> </ul>
keepalive	The interval at which messages are sent to the BGP neighbor to keep the mutual BGP session established. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
machines	The traffic managers that are to use this neighbor <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_bandwidth

URI Endpoint: /api/tm/8.1/config/active/bandwidth

A Bandwidth class, which can be assigned to a virtual server or pool in order to limit the number of bytes per second used by inbound or outbound traffic.

```
# Example usage of vtm_bandwidth resource...
resource "vtm_bandwidth" "my_bandwidth" {
  name = "MyBandwidth"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
maximum	The maximum bandwidth to allocate to connections that are associated with this bandwidth class (in kbits/second). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10000"</li> </ul>



Property	Description
note	A description of this bandwidth class. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sharing	The scope of the bandwidth class. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "cluster"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"cluster": Bandwidth is shared across all traffic managers</li> <li>"connection": Each connection can use the maximum rate</li> <li>"machine": Bandwidth is shared per traffic manager</li> </ul> </li> </ul>

### vtm\_cloud\_api\_credential

URI Endpoint: /api/tm/8.1/config/active/cloud\_api\_credentials

Cloud credentials used in cloud API calls

```
# Example usage of vtm_cloud_api_credential resource...
resource "vtm_cloud_api_credential" "my_cloud_api_credential" {
  name = "MyCloudApiCredential"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
api_server	The vCenter server hostname or IP address. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cloud_api_timeout	The traffic manager creates and destroys nodes via API calls. This setting specifies (in seconds) how long to wait for such calls to complete. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "200"</li> </ul>
cred1	The first part of the credentials for the cloud user. Typically this is some variation on the username concept. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
cred2	The second part of the credentials for the cloud user. Typically this is some variation on the password concept. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cred3	The third part of the credentials for the cloud user. Typically this is some variation on the authentication token concept. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
script	The script to call for communication with the cloud API. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
update_interval	The traffic manager will periodically check the status of the cloud through an API call. This setting specifies the interval between such updates. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>

### vtm\_custom

URI Endpoint: /api/tm/8.1/config/active/custom

Custom configuration sets store arbitrary named values. These values can be read by SOAP or REST clients.

```
# Example usage of vtm_custom resource...
resource "vtm_custom" "my_custom" {
  name = "MyCustom"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

Property	Description
string_lists	This table contains named lists of strings <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): Name of list (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>value (List(String)): Named list of user-specified strings. (Required)</li> </ul> </li> </ul>
string_lists_json	A Traffic Manager REST-compatible JSON representation of the "string_lists" table property. Use this field with the "vtm_custom_string_lists_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_dns\_server\_zone

URI Endpoint: /api/tm/8.1/config/active/dns\_server/zones

The conf/dnsserver/zones/ file contains zone metadata

```
# Example usage of vtm_dns_server_zone resource...
resource "vtm_dns_server_zone" "my_dns_server_zone" {
  name = "MyDnsServerZone"
  origin =
  zonefile =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
origin	The domain origin of this Zone. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
zonefile	The Zone File encapsulated by this Zone. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_dns\_server\_zone\_file

URI Endpoint: /api/tm/8.1/config/active/dns\_server/zone\_files

The `conf/dnsserver/zonefiles/` directory contains files that define DNS zones.

```
# Example usage of vtm_dns_server_zone_file resource...
resource "vtm_dns_server_zone_file" "my_dns_server_zone_file" {
  content = ""
  name = "MyDnsServerZoneFile"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_event\_type

URI Endpoint: `/api/tm/8.1/config/active/event_types`

Configuration that ties actions to a set of events that trigger them.

```
# Example usage of vtm_event_type resource...
resource "vtm_event_type" "my_event_type" {
  name = "MyEventType"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
actions	The actions triggered by events matching this event type, as a list of action references. <ul style="list-style-type: none"> <li>Type: List(Reference(config-event-action))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
built_in	If set to Yes this indicates that this configuration is built-in (provided as part of the software) and must not be deleted or edited. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
cloudcredentials_event_tags	Cloud credentials event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
cloudcredentials_objects	Cloud credentials object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
config_event_tags	Configuration file event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
faulttolerance_event_tags	Fault tolerance event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
general_event_tags	General event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
glb_event_tags	GLB service event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
glb_objects	GLB service object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
java_event_tags	Java event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
licensekeys_event_tags	License key event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
licensekeys_objects	License key object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
locations_event_tags	Location event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
locations_objects	Location object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
monitors_event_tags	Monitor event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
monitors_objects	Monitors object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	A description of this event type. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
pools_event_tags	Pool key event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
pools_objects	Pool object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
protection_event_tags	Service protection class event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
protection_objects	Service protection class object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rules_event_tags	Rule event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rules_objects	Rule object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
slm_event_tags	SLM class event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
slm_objects	SLM class object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_event_tags	SSL event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sslhw_event_tags	SSL hardware event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trafficscript_event_tags	TrafficScript event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
vservers_event_tags	Virtual server event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
vservers_objects	Virtual server object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
zxtms_event_tags	Traffic manager event tags <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
zxtms_objects	Traffic manager object names <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_extra\_file

URI Endpoint: /api/tm/8.1/config/active/extra\_files

A user-uploaded file. Such files can be used in TrafficScript code using the `resource.get` function.

```
# Example usage of vtm_extra_file resource...
resource "vtm_extra_file" "my_extra_file" {
  content = ""
  name = "MyExtraFile"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_glb\_service

URI Endpoint: /api/tm/8.1/config/active/glb\_services

A global load balancing service is used by a virtual server to modify DNS requests in order load balance data across different GLB locations.

```
# Example usage of vtm_glb_service resource...
resource "vtm_glb_service" "my_glb_service" {
  name = "MyGlbService"
}
```



Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>• Type: string</li> <li>• Required: true</li> </ul>
algorithm	Defines the global load balancing algorithm to be used. <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "hybrid"</li> <li>• Permitted values:               <ul style="list-style-type: none"> <li>"chained": Sends traffic to one location at a time, until that location fails where the next one in the chain is used.</li> <li>"geo": Distributes traffic based solely on the geographic location of each client.</li> <li>"hybrid": Distribute traffic based on both the load and geographic location.</li> <li>"load": Distributes traffic based on the current load to each location.</li> <li>"round_robin": Distributes traffic by assigning each request to a new location in turn. Over a period of time, all locations will receive the same number of requests.</li> <li>"weighted_random": Distributes traffic in a random way, but according to a weighted policy defined by individual location weights</li> </ul> </li> </ul>
all_monitors_needed	Do all monitors assigned to a location need to report success in order for it to be considered healthy? <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
autorecovery	The last location to fail will be available as soon as it recovers. <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
chained_auto_failback	Enable/Disable automatic failback mode. <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
chained_location_order	The locations this service operates for and defines the order in which locations fail. <ul style="list-style-type: none"> <li>• Type: List(String)</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
disable_on_failure	Locations recovering from a failure will become disabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
dnssec_keys	A table mapping domains to the private keys that authenticate them <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>domain (String): A domain authenticated by the associated private keys. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>ssl_key (Set(String)): Private keys that authenticate the associated domain. (Required)</li> </ul> </li> </ul>
dnssec_keys_json	A Traffic Manager REST-compatible JSON representation of the "dnssec_keys" table property. Use this field with the "vtm_glb_service_dnssec_keys_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
domains	The domains shown here should be a list of Fully Qualified Domain Names that you would like to balance globally. Responses from the back end DNS servers for queries that do not match this list will be forwarded to the client unmodified. Note: "*" may be used as a wild card. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
enabled	Enable/Disable our response manipulation of DNS. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
geo_effect	How much should the locality of visitors affect the choice of location used? This value is a percentage, 0% means that no locality information will be used, and 100% means that locality will always control which location is used. Values between the two extremes will act accordingly. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "50"</li> </ul>
last_resort_response	The response to be sent in case there are no locations available. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
location_draining	<p>This is the list of locations for which this service is draining. A location that is draining will never serve any of its service IP addresses for this domain. This can be used to take a location off-line.</p> <ul style="list-style-type: none"> <li>• Type: Set(String)</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
location_settings	<p>Table containing location specific settings.</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- location (String): Location to which the associated settings apply. (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- weight (UInt): Weight for this location, for use by the weighted random algorithm.</li> <li>- ips (Set(String)): The IP addresses that are present in a location. If the Global Load Balancer decides to direct a DNS query to this location, then it will filter out all IPs that are not in this list. (Required)</li> <li>- monitors (Set(String)): The monitors that are present in a location.</li> </ul> </li> </ul>
location_settings_json	<p>A Traffic Manager REST-compatible JSON representation of the "location_settings" table property. Use this field with the "vtm_glb_service_location_settings_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
log_enabled	<p>Log connections to this GLB service?</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
log_filename	<p>The filename the verbose query information should be logged to. Appliances will ignore this.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: "%zeushome%/zxtm/log/services/%g.log"</li> </ul>
log_format	<p>The format of the log lines.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: "%t, %s, %l, %q, %g, %n, %d, %a"</li> </ul>
return_ips_on_fail	<p>Return all or none of the IPs under complete failure.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>

Property	Description
rules	Response rules to be applied in the context of the service, in order, comma separated. <ul style="list-style-type: none"> <li>Type: List(Reference(config-trafficscript))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
tll	The TTL for the DNS resource records handled by the GLB service. <ul style="list-style-type: none"> <li>Type: Int</li> <li>Required: false</li> <li>Default value: "-1"</li> </ul>

### vtm\_global\_settings

URI Endpoint: /api/tm/8.1/config/active/global\_settings

General settings that apply to every machine in the cluster.

```
# Example usage of vtm_global_settings resource...
resource "vtm_global_settings" "my_global_settings" {

}
```

Property	Description
accepting_delay	How often, in milliseconds, each traffic manager child process (that isn't listening for new connections) checks to see whether it should start listening for new connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "50"</li> </ul>
admin_honor_fallback_scsv	Whether or not the admin server, the internal control port and the config daemon honor the Fallback SCSV to protect connections against downgrade attacks. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
admin_ssl3_allow_rehandshake	Whether or not SSL3/TLS re-handshakes should be supported for admin server and internal connections. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "rfc5746"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"always": Always allow</li> <li>"never": Never allow</li> <li>"rfc5746": Only if client uses RFC 5746 (Secure Renegotiation Extension)</li> <li>"safe": Allow safe re-handshakes</li> </ul> </li> </ul>

Property	Description
admin_ssl3_ciphers	<p>The SSL ciphers to use for admin server and internal connections. For information on supported ciphers see the online help.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
admin_ssl3_diffie_hellman_key_length	<p>The length in bits of the Diffie-Hellman key for ciphers that use Diffie-Hellman key agreement for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "dh_2048"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"dh_1024": Use 1024 bit keys for Diffie-Hellman ciphers.</li> <li>"dh_2048": Use 2048 bit keys for Diffie-Hellman ciphers.</li> <li>"dh_3072": Use 3072 bit keys for Diffie-Hellman ciphers.</li> <li>"dh_4096": Use 4096 bit keys for Diffie-Hellman ciphers.</li> </ul> </li> </ul>
admin_ssl3_min_rehandshake_interval	<p>If SSL3/TLS re-handshakes are supported on the admin server, this defines the minimum time interval (in milliseconds) between handshakes on a single SSL3/TLS connection that is permitted. To disable the minimum interval for handshakes the key should be set to the value 0.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1000"</li> </ul>
admin_ssl_elliptic_curves	<p>The SSL elliptic curve preference list for admin and internal connections. The named curves P256, P384 and P521 may be configured.</p> <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
admin_ssl_insert_extra_fragment	<p>Whether or not SSL3 and TLS1 use one-byte fragments as a BEAST countermeasure for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
admin_ssl_max_handshake_message_size	<p>The maximum size (in bytes) of SSL handshake messages that the admin server and internal connections will accept. To accept any size of handshake message the key should be set to the value 0.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10240"</li> </ul>

Property	Description
admin_ssl_signature_algorithms	<p>The SSL signature algorithms preference list for admin and internal connections using TLS version 1.2 or higher. For information on supported algorithms see the online help.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
admin_support_ssl3	<p>Whether or not SSL3 support is enabled for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
admin_support_tls1	<p>Whether or not TLS1.0 support is enabled for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
admin_support_tls1_1	<p>Whether or not TLS1.1 support is enabled for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
admin_support_tls1_2	<p>Whether or not TLS1.2 support is enabled for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
admin_support_tls1_3	<p>Whether or not TLS1.3 support is enabled for admin server and internal connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
afm_enabled	<p>Is the application firewall enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
appliance_bootloader_password	<p>The password used to protect the bootloader. An empty string means there will be no protection.</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
appliance_return_path_routing_enabled	<p>Whether or not the traffic manager will attempt to route response packets back to clients via the same route on which the corresponding request arrived. Note that this applies only to the last hop of the route - the behaviour of upstream routers cannot be altered by the traffic manager.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
optimizer_max_dependent_fetch_size	<p>The maximum size of a dependent resource that can undergo Web Accelerator optimization. Any content larger than this size will not be optimized. Units of KB and MB can be used, no postfix denotes bytes. A value of 0 disables the limit.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "2MB"</li> </ul>
optimizer_max_original_content_buffer_size	<p>The maximum size of unoptimized content buffered in the traffic manager for a single backend response that is undergoing Web Accelerator optimization. Responses larger than this will not be optimized. Note that if the backend response is compressed then this setting pertains to the compressed size, before Web Accelerator decompresses it. Units of KB and MB can be used, no postfix denotes bytes. Value range is 1 - 128MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "2MB"</li> </ul>
optimizer_watchdog_interval	<p>The period of time (in seconds) after which a previous failure will no longer count towards the watchdog limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "300"</li> </ul>
optimizer_watchdog_limit	<p>The maximum number of times the Web Accelerator sub-process will be started or restarted within the interval defined by the optimizer_watchdog_interval setting. If the process fails this many times, it must be restarted manually from the Diagnose page. Zero means no limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>
auditlog_via_eventd	<p>Whether to mirror the audit log to EventD.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
auditlog_via_syslog	<p>Whether to output audit log message to the syslog.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
auth_saml_key_lifetime	Lifetime in seconds of cryptographic keys used to decrypt SAML SP sessions stored externally (client-side). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "86400"</li> </ul>
auth_saml_key_rotation_interval	Rotation interval in seconds for cryptographic keys used to encrypt SAML SP sessions stored externally (client-side). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "14400"</li> </ul>
autoscaler_verbose	Whether or not detailed messages about the autoscaler's activity are written to the error log. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
bgp_as_number	The number of the BGP AS in which the traffic manager will operate. Must be entered in decimal. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65534"</li> </ul>
bgp_enabled	Whether BGP Route Health Injection is enabled <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
chunk_size	The default chunk size for reading/writing requests. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "16384"</li> </ul>
client_first_opt	Whether or not your traffic manager should make use of TCP optimisations to defer the processing of new client-first connections until the client has sent some data. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
cluster_comms_allow_update_default	The default value of allow_update for new cluster members. If you have cluster members joining from less trusted locations (such as cloud instances) this can be set to false in order to make them effectively "read-only" cluster members. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>



Property	Description
cluster_comms_allowed_update_hosts	<p>The hosts that can contact the internal administration port on each traffic manager. This should be a list containing IP addresses, CIDR IP subnets, and localhost; or it can be set to all to allow any host to connect.</p> <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: "all"</li> </ul>
cluster_comms_state_sync_interval	<p>How often to propagate the session persistence and bandwidth information to other traffic managers in the same cluster. Set this to 0 (zero) to disable propagation. Note that a cluster using "unicast" heartbeat messages cannot turn off these messages.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>
cluster_comms_state_sync_timeout	<p>The maximum amount of time to wait when propagating session persistence and bandwidth information to other traffic managers in the same cluster. Once this timeout is hit the transfer is aborted and a new connection created.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "6"</li> </ul>
cluster_identifier	<p>Cluster identifier. Generally supplied by Services Director.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_idle_connections_max	<p>The maximum number of unused HTTP keepalive connections with back-end nodes that the traffic manager should maintain for re-use. Setting this to 0 (zero) will cause the traffic manager to auto-size this parameter based on the available number of file-descriptors.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_idle_timeout	<p>How long an unused HTTP keepalive connection should be kept before it is discarded.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
connection_listen_queue_size	<p>The listen queue size for managing incoming connections. It may be necessary to increase the system's listen queue size if this value is altered. If the value is set to 0 then the default system setting will be used.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
connection_max_accepting	<p>Number of processes that should accept new connections. Only this many traffic manager child processes will listen for new connections at any one time. Setting this to 0 (zero) will cause your traffic manager to select an appropriate default value based on the architecture and number of CPUs.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
connection_multiple_accept	<p>Whether or not the traffic manager should try to read multiple new connections each time a new client connects. This can improve performance under some very specific conditions. However, in general it is recommended that this be set to 'false'.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
connection_udp_read_multiple	<p>Whether or not the traffic manager should try to read multiple UDP packets from clients each time the kernel reports data received from clients. This can improve performance for the situation with high UDP traffic throughput from clients to the traffic manager. Therefore, in general it is recommended that this be set to 'Yes'.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
dns_max_ttl	<p>Maximum Time To Live (expiry time) for entries in the DNS cache.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "86400"</li> </ul>
dns_min_ttl	<p>Minimum Time To Live (expiry time) for entries in the DNS cache.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "86400"</li> </ul>
dns_negative_expiry	<p>Expiry time for failed lookups in the DNS cache.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "60"</li> </ul>
dns_size	<p>Maximum number of entries in the DNS cache.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "10867"</li> </ul>
dns_timeout	<p>Timeout for receiving a response from a DNS server.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "12"</li> </ul>

Property	Description
ec2_access_key_id	<p>Deprecated: This key is unused. Amazon authentication credentials are now extracted from IAM Roles assigned to an EC2 instance.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ec2_awstool_timeout	<p>The maximum amount of time requests to the AWS Query API can take before timing out.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
ec2_metadata_server	<p>URL for the EC2 metadata server, <a href="http://169.254.169.254/latest/metadata">http://169.254.169.254/latest/metadata</a> for example.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ec2_query_server	<p>URL for the Amazon EC2 endpoint, <a href="https://ec2.amazonaws.com/">https://ec2.amazonaws.com/</a> for example.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ec2_secret_access_key	<p>Deprecated: This key is unused. Amazon authentication credentials are now extracted from IAM Roles assigned to an EC2 instance.</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ec2_verify_query_server_cert	<p>Whether to verify Amazon EC2 endpoint's certificate using CA(s) present in SSL Certificate Authorities Catalog.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
eventing_mail_interval	<p>The minimum length of time that must elapse between alert emails being sent. Where multiple alerts occur inside this timeframe, they will be retained and sent within a single email rather than separately.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
eventing_max_attempts	<p>The number of times to attempt to send an alert email before giving up.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>

Property	Description
fault_tolerance_arp_count	The number of ARP packets a traffic manager should send when an IP address is raised. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
fault_tolerance_auto_failback	Whether or not traffic IPs automatically move back to machines that have recovered from a failure and have dropped their traffic IPs. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
fault_tolerance_autofailback_delay	Configure the delay of automatic failback after a previous failover event. This setting has no effect if autofailback is disabled. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
fault_tolerance_child_timeout	How long the traffic manager should wait for status updates from any of the traffic manager's child processes before assuming one of them is no longer servicing traffic. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
fault_tolerance_frontend_check_ips	The IP addresses used to check front-end connectivity. The text %gateway% will be replaced with the default gateway on each system. Set this to an empty string if the traffic manager is on an Intranet with no external connectivity. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: "%gateway%"</li> </ul>
fault_tolerance_heartbeat_method	The method traffic managers should use to exchange cluster heartbeat messages. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "unicast"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"multicast": multicast</li> <li>"unicast": unicast</li> </ul> </li> </ul>
fault_tolerance_igmp_interval	The interval between unsolicited periodic IGMP Membership Report messages for Multi-Hosted Traffic IP Groups. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>

Property	Description
fault_tolerance_monitor_interval	The frequency, in milliseconds, that each traffic manager machine should check and announce its connectivity. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "500"</li> </ul>
fault_tolerance_monitor_timeout	How long, in seconds, each traffic manager should wait for a response from its connectivity tests or from other traffic manager machines before registering a failure. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
fault_tolerance_multicast_address	The multicast address and port to use to exchange cluster heartbeat messages. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "239.100.1.1:9090"</li> </ul>
fault_tolerance_unicast_port	The unicast UDP port to use to exchange cluster heartbeat messages. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "9090"</li> </ul>
fault_tolerance_use_bind_ip	Whether or not cluster heartbeat messages should only be sent and received over the management network. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
fault_tolerance_verbose	Whether or not a traffic manager should log all connectivity tests. This is very verbose, and should only be used for diagnostic purposes. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
fips_enabled	Enable FIPS Mode (requires software restart). <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ftp_data_bind_low	Whether or not the traffic manager should permit use of FTP data connection source ports lower than 1024. If No the traffic manager can completely drop root privileges, if Yes some or all privileges may be retained in order to bind to low ports. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
glb_verbose	Write a message to the logs for every DNS query that is load balanced, showing the source IP address and the chosen datacenter. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
historical_activity_keep_days	Number of days to store historical traffic information, if set to 0 the data will be kept indefinitely. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "90"</li> </ul>
ip_appliance_returnpath	A table of MAC to IP address mappings for each router where return path routing is required. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>mac (String): The MAC address of a router the software is connected to. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>ipv4 (String): The MAC address to IPv4 address mapping of a router the software is connected to. The * (asterisk) in the key name is the MAC address, the value is the IP address.</li> <li>ipv6 (String): The MAC address to IPv6 address mapping of a router the software is connected to. The * (asterisk) in the key name is the MAC address, the value is the IP address.</li> </ul> </li> </ul>
ip_appliance_returnpath_json	A Traffic Manager REST-compatible JSON representation of the "ip_appliance_returnpath" table property. Use this field with the "vtm_global_settings_appliance_returnpath_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
java_classpath	CLASSPATH to use when starting the Java runner. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
java_command	Java command to use when starting the Java runner, including any additional options. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "java -server"</li> </ul>

Property	Description
java_enabled	<p>Whether or not Java support should be enabled. If this is set to No, then your traffic manager will not start any Java processes. Java support is only required if you are using the TrafficScript java.run() function.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
java_lib	<p>Java library directory for additional jar files. The Java runner will load classes from any .jar files stored in this directory, as well as the * .jar files and classes stored in traffic manager's catalog.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
java_max_connections	<p>Maximum number of simultaneous Java requests. If there are more than this many requests, then further requests will be queued until the earlier requests are completed. This setting is per-CPU, so if your traffic manager is running on a machine with 4 CPU cores, then each core can make this many requests at one time.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "256"</li> </ul>
java_session_age	<p>Default time to keep a Java session.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "86400"</li> </ul>
kerberos_verbose	<p>Whether or not a traffic manager should log all Kerberos related activity. This is very verbose, and should only be used for diagnostic purposes.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
license_servers	<p>A list of license servers for FLA licensing. A license server should be specified as a &lt;ip/host&gt;:&lt;port&gt; pair.</p> <ul style="list-style-type: none"> <li>• Type: Set(String)</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
log_error_level	<p>The minimum severity of events/alerts that should be logged to disk. INFO will log all events; a higher severity setting will log fewer events. More fine-grained control can be achieved using events and actions.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "info"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"fatal": Only fatal errors are logged</li> <li>"info": All events are logged to disk</li> <li>"serious": Only serious errors or worse</li> <li>"warn": Only warnings and errors are logged</li> </ul> </li> </ul>
log_export_auth_hec_token	<p>The HTTP Event Collector token to use for HTTP authentication with a Splunk server.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
log_export_auth_http	<p>The HTTP authentication method to use when exporting log entries.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"basic": Basic (Username and Password)</li> <li>"none": None</li> <li>"splunk": Splunk (HEC token)</li> </ul> </li> </ul>
log_export_auth_password	<p>The password to use for HTTP basic authentication.</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
log_export_auth_username	<p>The username to use for HTTP basic authentication.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
log_export_enabled	<p>Monitor log files and export entries to the configured endpoint.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
log_export_endpoint	<p>The URL to which log entries should be sent. Entries are sent using HTTP(S) POST requests.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>



Property	Description
log_export_request_timeout	The number of seconds after which HTTP requests sent to the configured endpoint will be considered to have failed if no response is received. A value of 0 means that HTTP requests will not time out. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
log_export_tls_verify	Whether the server certificate should be verified when connecting to the endpoint. If enabled, server certificates that do not match the server name, are self-signed, have expired, have been revoked, or that are signed by an unknown CA will be rejected. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
log_flush_time	How long to wait before flushing the request log files for each virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
log_log_file	The file to log event messages to. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "%zeushome%/zxtm/log/errors"</li> </ul>
log_rate	The maximum number of connection errors logged per second when connection error reporting is enabled. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "50"</li> </ul>
log_reopen	How long to wait before re-opening request log files, this ensures that log files will be recreated in the case of log rotation. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
log_time	The minimum time between log messages for log intensive features such as SLM. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
max_fds	The maximum number of file descriptors that your traffic manager will allocate. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1048576"</li> </ul>

Property	Description
monitor_memory_size	<p>The maximum number of each of nodes, pools or locations that can be monitored. The memory used to store information about nodes, pools and locations is allocated at start-up, so the traffic manager must be restarted after changing this setting.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>
ospfv2_area	<p>The OSPF area in which the traffic manager will operate. May be entered in decimal or IPv4 address format.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "0.0.0.1"</li> </ul>
ospfv2_area_type	<p>The type of OSPF area in which the traffic manager will operate. This must be the same for all routers in the area, as required by OSPF.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "normal"</li> <li>Permitted values: "normal": Normal area "nssa": Not So Stubby Area (RFC3101) "stub": Stub area</li> </ul>
ospfv2_authentication_key_id_a	<p>OSPFv2 authentication key ID. If set to 0, which is the default value, the key is disabled.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ospfv2_authentication_key_id_b	<p>OSPFv2 authentication key ID. If set to 0, which is the default value, the key is disabled.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ospfv2_authentication_shared_secret_a	<p>OSPFv2 authentication shared secret (MD5). If set to blank, which is the default value, the key is disabled.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ospfv2_authentication_shared_secret_b	<p>OSPFv2 authentication shared secret (MD5). If set to blank, which is the default value, the key is disabled.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ospfv2_dead_interval	The number of seconds before declaring a silent router down. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "40"</li> </ul>
ospfv2_enabled	Whether OSPFv2 Route Health Injection is enabled <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ospfv2_hello_interval	The interval at which OSPF "hello" packets are sent to the network. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
protection_conncount_size	The amount of shared memory reserved for an inter-process table of combined connection counts, used by all Service Protection classes that have per_process_connection_count set to No. The amount is specified as an absolute size, eg 20MB. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "20MB"</li> </ul>
rate_class_limit	The maximum number of Rate classes that can be created. Approximately 100 bytes will be pre-allocated per Rate class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "25000"</li> </ul>
recent_connections_max_per_process	How many recently closed connections each traffic manager process should save. These saved connections will be shown alongside currently active connections when viewing the Connections page. You should set this value to 0 in a benchmarking or performance-critical environment. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "500"</li> </ul>
recent_connections_retain_time	The amount of time for which snapshots will be retained on the Connections page. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
recent_connections_snapshot_size	The maximum number of connections each traffic manager process should show when viewing a snapshot on the Connections page. This value includes both currently active connections and saved connections. If set to 0 all active and saved connection will be displayed on the Connections page. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "500"</li> </ul>

Property	Description
remote_licensing_comm_channel_enabled	Whether to create a Communications Channel agent to send and receive messages from the Services Director Registration Server. This will be disabled when performing self-registration with a Services Director which does not support this feature. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
remote_licensing_comm_channel_port	The port number the Services Director instance is using for access to the traffic manager Communications Channel. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "8102"</li> </ul>
remote_licensing_owner	The Owner of a Services Director instance, used for self-registration. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_owner_secret	The secret associated with the Owner. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_policy_id	The auto-accept Policy ID that this instance should attempt to use. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_registration_server	A Services Director address for self-registration. A registration server should be specified as a <ip/host>:<port> pair. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_server_certificate	The certificate of a Services Director instance, used for self-registration. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rest_api_auth_timeout	The length of time after a successful request that the authentication of a given username and password will be cached for an IP address. A setting of 0 disables the cache forcing every REST request to be authenticated which will adversely affect performance. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "120"</li> </ul>
rest_api_enabled	Whether or not the REST service is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
rest_api_http_max_header_length	The maximum allowed length in bytes of a HTTP request's headers. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>
rest_api_maxfds	Maximum number of file descriptors that the REST API will allocate. The REST API must be restarted for a change to this setting to take effect. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1048576"</li> </ul>
rest_api_replicate_absolute	Configuration changes will be replicated across the cluster after this period of time, regardless of whether additional API requests are being made. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "20"</li> </ul>
rest_api_replicate_lull	Configuration changes made via the REST API will be propagated across the cluster when no further API requests have been made for this period of time. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
rest_api_replicate_timeout	The period of time after which configuration replication across the cluster will be cancelled if it has not completed. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
security_login_banner	Banner text displayed on the Admin Server login page and before logging in to appliance SSH servers. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_login_banner_accept	Whether or not users must explicitly agree to the displayed login_banner text before logging in to the Admin Server. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
security_login_delay	The number of seconds before another login attempt can be made after a failed attempt. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4"</li> </ul>

Property	Description
security_max_login_attempts	The number of sequential failed login attempts that will cause a user account to be suspended. Setting this to 0 disables this feature. To apply this to users who have never successfully logged in, track_unknown_users must also be enabled. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_max_login_external	Whether or not usernames blocked due to the max_login_attempts limit should also be blocked from authentication against external services (such as LDAP and RADIUS). <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
security_max_login_suspension_time	The number of minutes to suspend users who have exceeded the max_login_attempts limit. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "15"</li> </ul>
security_password_allow_consecutive_chars	Whether or not to allow the same character to appear consecutively in passwords. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
security_password_changes_per_day	The maximum number of times a password can be changed in a 24-hour period. Set to 0 to disable this restriction. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_password_min_alpha_chars	Minimum number of alphabetic characters a password must contain. Set to 0 to disable this restriction. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_password_min_length	Minimum number of characters a password must contain. Set to 0 to disable this restriction. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_password_min_numeric_chars	Minimum number of numeric characters a password must contain. Set to 0 to disable this restriction. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
security_password_min_special_chars	<p>Minimum number of special (non-alphanumeric) characters a password must contain. Set to 0 to disable this restriction.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_password_min_uppercase_chars	<p>Minimum number of uppercase characters a password must contain. Set to 0 to disable this restriction.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_password_reuse_after	<p>The number of times a password must have been changed before it can be reused. Set to 0 to disable this restriction.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_post_login_banner	<p>Banner text to be displayed on the appliance console after login.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
security_track_unknown_users	<p>Whether to remember past login attempts from usernames that are not known to exist (should be set to false for an Admin Server accessible from the public Internet). This does not affect the audit log.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
security_ui_page_banner	<p>Banner text to be displayed on all Admin Server pages.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
session_asp_cache_size	<p>The maximum number of entries in the ASP session persistence cache. This is used for storing session mappings for ASP session persistence. Approximately 100 bytes will be pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "32768"</li> </ul>
session_ip_cache_expiry	<p>IP session persistence cache expiry time in seconds. A session will not be reused if the time since it was last used exceeds this value. 0 indicates no expiry timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
session_ip_cache_size	<p>The maximum number of entries in the IP session persistence cache. This is used to provide session persistence based on the source IP address. Approximately 100 bytes will be pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "32768"</li> </ul>
session_j2ee_cache_expiry	<p>J2EE session persistence cache expiry time in seconds. A session will not be reused if the time since it was last used exceeds this value. 0 indicates no expiry timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
session_j2ee_cache_size	<p>The maximum number of entries in the J2EE session persistence cache. This is used for storing session mappings for J2EE session persistence. Approximately 100 bytes will be pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "32768"</li> </ul>
session_ssl_cache_size	<p>The maximum number of entries in the SSL session persistence cache. This is used to provide session persistence based on the SSL session ID. Approximately 200 bytes will be pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "32768"</li> </ul>
session_universal_cache_expiry	<p>Universal session persistence cache expiry time in seconds. A session will not be reused if the time since it was last used exceeds this value. 0 indicates no expiry timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
session_universal_cache_size	<p>The maximum number of entries in the global universal session persistence cache. This is used for storing session mappings for universal session persistence. Approximately 100 bytes will be pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "32768"</li> </ul>
shared_pool_size	<p>The size of the shared memory pool used for shared storage across worker processes (e.g. bandwidth shared data). This is specified as either a percentage of system RAM, 5% for example, or an absolute size such as 10MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "10MB"</li> </ul>



Property	Description
slm_class_limit	<p>The maximum number of SLM classes that can be created. Approximately 100 bytes will be pre-allocated per SLM class.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1024"</li> </ul>
snmp_user_counters	<p>The number of user defined SNMP counters. Approximately 100 bytes will be pre-allocated at start-up per user defined SNMP counter.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
so_rbuff_size	<p>The size of the operating system's read buffer. A value of 0 (zero) means to use the OS default; in normal circumstances this is what should be used.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
so_wbuff_size	<p>The size of the operating system's write buffer. A value of 0 (zero) means to use the OS default; in normal circumstances this is what should be used.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
soap_idle_minutes	<p>The number of minutes that the SOAP server should remain idle before exiting. The SOAP server has a short startup delay the first time a SOAP request is made, subsequent SOAP requests don't have this delay.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
socket_optimizations	<p>Whether or not the traffic manager should use potential network socket optimisations. If set to auto, a decision will be made based on the host platform.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "auto"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"auto": Decide based on local platform</li> <li>"no": Disable socket optimizations</li> <li>"yes": Enable socket optimizations</li> </ul> </li> </ul>

Property	Description
ssl_allow_rehandshake	<p>Whether or not SSL/TLS re-handshakes should be supported. Enabling support for re-handshakes can expose services to Man-in-the-Middle attacks. It is recommended that only "safe" handshakes be permitted, or none at all.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "safe"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"always": Always allow</li> <li>"never": Never allow</li> <li>"rfc5746": Only if client uses RFC 5746 (Secure Renegotiation Extension)</li> <li>"safe": Allow safe re-handshakes</li> </ul> </li> </ul>
ssl_cache_enabled	<p>Whether or not the SSL server session cache is enabled, unless overridden by virtual server settings.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_cache_expiry	<p>How long the SSL session IDs for SSL decryption should be stored for.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1800"</li> </ul>
ssl_cache_per_virtualserver	<p>Whether an SSL session created by a given virtual server can only be resumed by a connection to the same virtual server.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_cache_size	<p>How many entries the SSL session ID cache should hold. This cache is used to cache SSL sessions to help speed up SSL handshakes when performing SSL decryption. Each entry will allocate approximately 1.5kB of metadata.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "6151"</li> </ul>
ssl_cipher_suites	<p>The SSL/TLS cipher suites preference list for SSL/TLS connections, unless overridden by virtual server or pool settings. For information on supported cipher suites see the online help.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ssl_client_cache_enabled	<p>Whether or the SSL client cache will be used, unless overridden by pool settings.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_client_cache_expiry	<p>How long in seconds SSL sessions should be stored in the client cache for, by default. Servers returning session tickets may also provide a lifetime hint, which will be used if it is less than this value.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "14400"</li> </ul>
ssl_client_cache_size	<p>How many entries the SSL client session cache should hold, per child. This cache is used to cache SSL sessions to help speed up SSL handshakes when performing SSL encryption. Each entry will require approx 100 bytes of memory plus space for either an SSL session id or an SSL session ticket, which may be as small as 16 bytes or may be as large as a few kilobytes, depending upon the server behavior.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1024"</li> </ul>
ssl_client_cache_tickets_enabled	<p>Whether or not session tickets, including TLS &gt;= 1.3 PSKs, may be requested and stored in the SSL client cache.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_crl_mem_size	<p>How much shared memory to allocate for loading Certificate Revocation Lists. This should be at least 3 times the total size of all CRLs on disk. This is specified as either a percentage of system RAM, 1% for example, or an absolute size such as 10MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "5MB"</li> </ul>
ssl_diffie_hellman_modulus_size	<p>The size in bits of the modulus for the domain parameters used for cipher suites that use finite field Diffie-Hellman key agreement.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "dh_2048"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"dh_1024": 1024 bit modulus</li> <li>"dh_2048": 2048 bit modulus</li> <li>"dh_3072": 3072 bit modulus</li> <li>"dh_4096": 4096 bit modulus</li> </ul> </li> </ul>

Property	Description
ssl_elliptic_curves	The SSL/TLS elliptic curve preference list for SSL/TLS connections using TLS version 1.0 or higher, unless overridden by virtual server or pool settings. For information on supported curves see the online help. <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_accel	Whether or not the SSL hardware is an "accelerator" (faster than software). By default the traffic manager will only use the SSL hardware if a key requires it (i.e. the key is stored on secure hardware and the traffic manager only has a placeholder/identifier key). With this option enabled, your traffic manager will instead try to use hardware for all SSL decrypts. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_hardware_azure_client_id	The client identifier used when accessing the Microsoft Azure Key Vault. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_azure_client_secret	The client secret used when accessing the Microsoft Azure Key Vault. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_azure_vault_url	The URL for the REST API of the Microsoft Azure Key Vault. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_azure_verify_rest_api_cert	Whether or not the Azure Key Vault REST API certificate should be verified. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_hardware_driver_pkcs11_debug	Print verbose information about the PKCS11 hardware security module to the event log. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_hardware_driver_pkcs11_library	The location of the PKCS#11 library for your SSL hardware if it is not in a standard location. The traffic manager will search the standard locations by default. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ssl_hardware_driver_pkcs11_slot_desc	<p>The label of the SSL Hardware slot to use. Only required if you have multiple HW accelerator slots.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_driver_pkcs11_slot_type	<p>The type of SSL hardware slot to use.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "operator"</li> <li>Permitted values: "module": Module Protected "operator": Operator Card Set "softcard": Soft Card</li> </ul>
ssl_hardware_driver_pkcs11_user_pin	<p>The User PIN for the PKCS token (PKCS#11 devices only).</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_hardware_failure_count	<p>The number of consecutive failures from the SSL hardware that will be tolerated before the traffic manager assumes its session with the device is invalid and tries to log in again. This is necessary when the device reboots following a power failure.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
ssl_hardware_library	<p>The type of SSL hardware to use. The drivers for the SSL hardware should be installed and accessible to the traffic manager software.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: "azure": Microsoft Azure Key Vault "none": None "pkcs11": PKCS#11</li> </ul>
ssl_honor_fallback_scsv	<p>Whether or not ssl-decrypting Virtual Servers honor the Fallback SCSV to protect connections against downgrade attacks.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_insert_extra_fragment	<p>Whether or not SSL3 and TLS1 use one-byte fragments as a BEAST countermeasure.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
ssl_log_keys	<p>Whether SSL connection key logging should be available via the <code>ssl.sslkeylogline()</code> TrafficScript function. If this setting is disabled then <code>ssl.sslkeylogline()</code> will always return the empty string.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_max_handshake_message_size	<p>The maximum size (in bytes) of SSL handshake messages that SSL connections will accept. To accept any size of handshake message the key should be set to the value 0.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10240"</li> </ul>
ssl_middlebox_compatibility	<p>Whether or not TLS 1.3 middlebox compatibility mode as described in RFC 8446 appendix D.4 will be used in connections to pool nodes, unless overridden by pool settings.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_min_rehandshake_interval	<p>If SSL3/TLS re-handshakes are supported, this defines the minimum time interval (in milliseconds) between handshakes on a single SSL3/TLS connection that is permitted. To disable the minimum interval for handshakes the key should be set to the value 0.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1000"</li> </ul>
ssl_ocsp_cache_size	<p>The maximum number of cached client certificate OCSP results stored. This cache is used to speed up OCSP checks against client certificates by caching results. Approximately 1040 bytes are pre-allocated per entry.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2048"</li> </ul>
ssl_ocsp_stapling_default_refresh_interval	<p>How long to wait before refreshing requests on behalf of the store of certificate status responses used by OCSP stapling, if we don't have an up-to-date OCSP response.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
ssl_ocsp_stapling_maximum_refresh_interval	<p>Maximum time to wait before refreshing requests on behalf of the store of certificate status responses used by OCSP stapling. (0 means no maximum.)</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "864000"</li> </ul>

Property	Description
ssl_ocsp_stapling_mem_size	<p>How much shared memory to allocate for the store of certificate status responses for OCSP stapling. This should be at least 2kB times the number of certificates configured to use OCSP stapling. This is specified as either a percentage of system RAM, 1% for example, or an absolute size such as 10MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "1MB"</li> </ul>
ssl_ocsp_stapling_time_tolerance	<p>How many seconds to allow the current time to be outside the validity time of an OCSP response before considering it invalid.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
ssl_ocsp_stapling_verify_response	<p>Whether the OCSP response signature should be verified before the OCSP response is cached.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_signature_algorithms	<p>The SSL/TLS signature algorithms preference list for SSL/TLS connections using TLS version 1.2 or higher, unless overridden by virtual server or pool settings. For information on supported algorithms see the online help.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_support_ssl3	<p>Whether or not SSL3 support is enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_support_tls1	<p>Whether or not TLS1.0 support is enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_support_tls1_1	<p>Whether or not TLS1.1 support is enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_support_tls1_2	<p>Whether or not TLS1.2 support is enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
ssl_support_tls1_3	Whether or not TLS1.3 support is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_tickets_enabled	Whether or not session tickets will be issued to and accepted from clients that support them, unless overridden by virtual server settings. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_tickets_reissue_policy	When an SSL session ticket will be reissued (ie when a new ticket will be generated for the same SSL session). <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "never"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"always": always</li> <li>"never": never</li> </ul> </li> </ul>
ssl_tickets_ticket_expiry	The length of time for which an SSL session ticket will be accepted by a virtual server after the ticket is created. If a ticket is reissued (if ssl!tickets!reissue_policy is set to 'always') this time starts at the time when the ticket was reissued. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "14400"</li> </ul>
ssl_tickets_ticket_key_expiry	The length of time for which an auto-generated SSL ticket key will be used to decrypt old session ticket, before being deleted from memory. This setting is ignored if there are any entries in the (REST-only) SSL ticket keys catalog. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "86400"</li> </ul>
ssl_tickets_ticket_key_rotation	The length of time for which an auto-generated SSL ticket key will be used to encrypt new session tickets, before a new SSL ticket key is generated. The ticket encryption key will be held in memory for ssl!tickets!ticket_key_expiry, so that tickets encrypted using the key can still be decrypted and used. This setting is ignored if there are any entries in the (REST-only) SSL ticket keys catalog. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "14400"</li> </ul>
ssl_tickets_time_tolerance	How many seconds to allow the current time to be outside the validity time of an SSL ticket before considering it invalid. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>



Property	Description
ssl_validate_server_certificates_catalog	Whether the traffic manager should validate that SSL server certificates form a matching key pair before the certificate gets used on an SSL decrypting virtual server. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
telemetry_enabled	Allow the reporting of anonymized usage data for product improvement and customer support purposes. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
tip_class_limit	The maximum number of Traffic IP Groups that can be created. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10000"</li> </ul>
trafficscript_data_local_size	The maximum amount of memory available to store TrafficScript data.local.set() information. This can be specified as a percentage of system RAM, 5% for example; or an absolute size such as 200MB. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "5%"</li> </ul>
trafficscript_data_size	The maximum amount of memory available to store TrafficScript data.set() information. This can be specified as a percentage of system RAM, 5% for example; or an absolute size such as 200MB. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "5%"</li> </ul>
trafficscript_execution_time_warning	Raise an event if a TrafficScript rule runs for more than this number of milliseconds in a single invocation. If you get such events repeatedly, you may want to consider re-working some of your TrafficScript rules. A value of 0 means no warnings will be issued. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "500"</li> </ul>
trafficscript_max_instr	The maximum number of instructions a TrafficScript rule will run. A rule will be aborted if it runs more than this number of instructions without yielding, preventing infinite loops. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "100000"</li> </ul>

Property	Description
trafficscript_memory_warning	<p>Raise an event if a TrafficScript rule requires more than this amount of buffered network data. If you get such events repeatedly, you may want to consider re-working some of your TrafficScript rules to use less memory or to stream the data that they process rather than storing it all in memory. This setting also limits the amount of data that can be returned by request.GetLine().</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1048576"</li> </ul>
trafficscript_regex_cache_size	<p>The maximum number of regular expressions to cache in TrafficScript. Regular expressions will be compiled in order to speed up their use in the future.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "57"</li> </ul>
trafficscript_regex_match_limit	<p>The maximum number of ways TrafficScript will attempt to match a regular expression at each position in the subject string, before it aborts the rule and reports a TrafficScript error.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10000000"</li> </ul>
trafficscript_regex_match_warning_percentage	<p>The percentage of regex_match_limit at which TrafficScript reports a performance warning.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
trafficscript_variable_pool_use	<p>Allow the pool.use and pool.select TrafficScript functions to accept variables instead of requiring literal strings. Enabling this feature has the following effects<sup>1</sup>. Your traffic manager may no longer be able to know whether a pool is in use.<sup>2</sup> Errors for pools that aren't in use will not be hidden.<sup>3</sup> Some settings displayed for a Pool may not be appropriate for the type of traffic being managed.<sup>4</sup> Pool usage information on the pool edit pages and config summary may not be accurate.<sup>5</sup> Monitors will run for all pools (with this option disabled monitors will only run for Pools that are used).</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
transaction_export_enabled	<p>Export metadata about transactions processed by the traffic manager to an external location.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
transaction_export_endpoint	<p>The endpoint to which transaction metadata should be exported. The endpoint is specified as a hostname or IP address with a port.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
transaction_export_tls	<p>Whether the connection to the specified endpoint should be encrypted.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
transaction_export_tls_verify	<p>Whether the server certificate presented by the endpoint should be verified, preventing a connection from being established if the certificate does not match the server name, is self-signed, is expired, is revoked, or has an unknown CA.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
watchdog_timeout	<p>The maximum time in seconds a process can fail to update its heartbeat, before the watchdog considers it to have stalled.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
web_cache_avg_path_length	<p>The estimated average length of the path (including query string) for resources being cached. An amount of memory equal to this figure multiplied by max_file_num will be allocated for storing the paths for cache entries. This setting can be increased if your web site makes extensive use of long URLs.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "512"</li> </ul>
web_cache_disk	<p>Whether or not to use a disk-backed (typically SSD) cache. If set to Yes cached web pages will be stored in a file on disk. This enables the traffic manager to use a cache that is larger than available RAM. The size setting should also be adjusted to select a suitable maximum size based on your disk space. Note that the disk caching is optimized for use with SSD storage.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
web_cache_disk_dir	<p>If disk caching is enabled, this sets the directory where the disk cache file will be stored. The traffic manager will create a file called webcache.data in this location. Note that the disk caching is optimized for use with SSD storage.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "%zeushome%/zxtm/internal"</li> </ul>

Property	Description
web_cache_max_file_num	<p>Maximum number of entries in the cache. Approximately 0.9 KB will be pre-allocated per entry for metadata, this is in addition to the memory reserved for the content cache and for storing the paths of the cached resources.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10000"</li> </ul>
web_cache_max_file_size	<p>Largest size of a cacheable object in the cache. This is specified as either a percentage of the total cache size, 2% for example, or an absolute size such as 20MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "2%"</li> </ul>
web_cache_max_path_length	<p>The maximum length of the path (including query string) for the resource being cached. If the path exceeds this length then it will not be added to the cache.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2048"</li> </ul>
web_cache_normalize_query	<p>Enable normalization (lexical ordering of the parameter-assignments) of the query string.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
web_cache_size	<p>The maximum size of the HTTP web page cache. This is specified as either a percentage of system RAM, 20% for example, or an absolute size such as 200MB.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "20%"</li> </ul>
web_cache_verbose	<p>Add an X-Cache-Info header to every HTTP response, showing whether the request and/or the response was cacheable.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

## vtm\_kerberos\_krb5conf

URI Endpoint: /api/tm/8.1/config/active/kerberos/krb5confs

A Kerberos krb5.conf file that provides the raw configuration for a Kerberos principal.

```
# Example usage of vtm_kerberos_krb5conf resource...
resource "vtm_kerberos_krb5conf" "my_kerberos_krb5conf" {
  content = ""
  name = "MyKerberosKrb5Conf"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_kerberos\_keytab

URI Endpoint: /api/tm/8.1/config/active/kerberos/keytabs

A Kerberos keytab file contains credentials to authenticate as (a number of) Kerberos principals.

```
# Example usage of vtm_kerberos_keytab resource...
resource "vtm_kerberos_keytab" "my_kerberos_keytab" {
  content = ""
  name = "MyKerberosKeytab"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_kerberos\_principal

URI Endpoint: /api/tm/8.1/config/active/kerberos/principals

A Kerberos principal can be used by the traffic manager to participate in a Kerberos realm.

```
# Example usage of vtm_kerberos_principal resource...
resource "vtm_kerberos_principal" "my_kerberos_principal" {
  keytab =
  name = "MyKerberosPrincipal"
  service =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
kdc	A list of <hostname/ip>:<port> pairs for Kerberos key distribution center (KDC) services to be explicitly used for the realm of the principal. If no KDCs are explicitly configured, DNS will be used to discover the KDC(s) to use. <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
keytab	The name of the Kerberos keytab file containing suitable credentials to authenticate as the specified Kerberos principal. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
krb5conf	The name of an optional Kerberos configuration file (krb5.conf). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
realm	The Kerberos realm where the principal belongs. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
service	The service name part of the Kerberos principal name the traffic manager should use to authenticate itself. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_license\_key

URI Endpoint: /api/tm/8.1/config/active/license\_keys

A license key is an encoded text file that controls what functionality is available from each traffic manager in the cluster. Every production traffic manager must have a valid licence key in order to function; a traffic manager without a license will operate as Community Edition, which provides most of the functionality, but places restrictions on bandwidth and cluster size.

```
# Example usage of vtm_license_key resource...
resource "vtm_license_key" "my_license_key" {
  content = ""
  name = "MyLicenseKey"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

## vtm\_location

URI Endpoint: /api/tm/8.1/config/active/locations

These are geographic locations as used by Global Load Balancing services. Such a location may not necessarily contain a traffic manager; instead it could refer to the location of a remote datacenter.

```
# Example usage of vtm_location resource...
resource "vtm_location" "my_location" {
  id =
  name = "MyLocation"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
identifier	The identifier of this location. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
latitude	The latitude of this location. <ul style="list-style-type: none"> <li>Type: Float</li> <li>Required: false</li> <li>Default value: "0.0"</li> </ul>
longitude	The longitude of this location. <ul style="list-style-type: none"> <li>Type: Float</li> <li>Required: false</li> <li>Default value: "0.0"</li> </ul>

Property	Description
note	A note, used to describe this location. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
type	Does this location contain traffic managers and configuration or is it a recipient of GLB requests? <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "config"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"config": Configuration</li> <li>"glb": GLB</li> </ul> </li> </ul>

### vtm\_log\_export

URI Endpoint: /api/tm/8.1/config/active/log\_export

Definitions of log files which should be exported to the analytics engine

```
# Example usage of vtm_log_export resource...
resource "vtm_log_export" "my_log_export" {
  name = "MyLogExport"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
appliance_only	Whether entries from the specified log files should be exported only from appliances. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
enabled	Export entries from the log files included in this category. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
files	The set of files to export as part of this category, specified as a list of glob patterns. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>



Property	Description
history	<p>How much historic log activity should be exported.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"all": Export all historic entries</li> <li>"none": Do not export any historic entries</li> <li>"recent": Export recent historic entries, according to the 'history_period' setting</li> </ul> </li> </ul>
history_period	<p>The number of days of historic log entries that should be exported.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
metadata	<p>This is table 'metadata'</p> <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): The name of a metadata item which should be sent to the analytics engine along with entries from these log files. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>value (String): Additional metadata to include with the log entries when exporting them to the configured endpoint. Metadata can be used by the system that is receiving the exported data to categorise and parse the log entries. (Required)</li> </ul> </li> </ul>
metadata_json	<p>A Traffic Manager REST-compatible JSON representation of the "metadata" table property. Use this field with the "vtm_log_export_metadata_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	<p>A description of this category of log files.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_monitor

URI Endpoint: /api/tm/8.1/config/active/monitors

Monitors check important remote services are running, by periodically sending them traffic and checking the response is correct. They are used by virtual servers to detect the failure of backend nodes.

```
# Example usage of vtm_monitor resource...
resource "vtm_monitor" "my_monitor" {
  name = "MyMonitor"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
back_off	Should the monitor slowly increase the delay after it has failed? <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
delay	The minimum time between calls to a monitor. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>
failures	The number of times in a row that a node must fail execution of the monitor before it is classed as unavailable. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>
health_only	Should this monitor only report health (ignore load)? <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
http_authentication	The HTTP basic-auth <user>:<password> to use for the test HTTP request. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_body_regex	A regular expression that the HTTP response body must match. If the response body content doesn't matter then set this to .* (match anything). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_host_header	The host header to use in the test HTTP request. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
http_path	The path to use in the test HTTP request. This must be a string beginning with a / (forward slash). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "/"</li> </ul>
http_status_regex	A regular expression that the HTTP status code must match. If the status code doesn't matter then set this to .* (match anything). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "^([234][0-9][0-9])\$"</li> </ul>
machine	The machine to monitor, where relevant this should be in the form <hostname>:<port>, for "ping" monitors the :<port> part must not be specified. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	A description of the monitor. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rtsp_body_regex	The regular expression that the RTSP response body must match. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rtsp_path	The path to use in the RTSP request (some servers will return 500 Internal Server Error unless this is a valid media file). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "/"</li> </ul>
rtsp_status_regex	The regular expression that the RTSP response status code must match. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "^([234][0-9][0-9])\$"</li> </ul>
scope	A monitor can either monitor each node in the pool separately and disable an individual node if it fails, or it can monitor a specific machine and disable the entire pool if that machine fails. GLB location monitors must monitor a specific machine. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "pernode"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"pernode": Node: Monitor each node in the pool separately</li> <li>"poolwide": Pool/GLB: Monitor a specified machine</li> </ul> </li> </ul>

Property	Description
script_arguments	<p>A table containing arguments and argument values to be passed to the monitor program.</p> <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): The name of the argument to be passed to the monitor program. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>value (String): The value of the argument to be passed to the monitor program. (Required)</li> <li>description (String): A description for the argument provided to the program.</li> </ul> </li> </ul>
script_arguments_json	<p>A Traffic Manager REST-compatible JSON representation of the "script_arguments" table property. Use this field with the "vtm_monitor_arguments_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
script_program	<p>The program to run. This must be an executable file, either within the monitor scripts directory or specified as an absolute path to some other location on the filesystem.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sip_body_regex	<p>The regular expression that the SIP response body must match.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sip_status_regex	<p>The regular expression that the SIP response status code must match.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "^([234][0-9][0-9])\$" </li> </ul>
sip_transport	<p>Which transport protocol the SIP monitor will use to query the server.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "udp"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"tcp": TCP</li> <li>"udp": UDP</li> </ul> </li> </ul>
tcp_close_string	<p>An optional string to write to the server before closing the connection.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
tcp_max_response_len	The maximum amount of data to read back from a server, use 0 for unlimited. Applies to TCP and HTTP monitors. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2048"</li> </ul>
tcp_response_regex	A regular expression to match against the response from the server. Applies to TCP monitors only. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: ".+"</li> </ul>
tcp_write_string	The string to write down the TCP connection. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
timeout	The maximum runtime for an individual instance of the monitor. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>
type	The internal monitor implementation of this monitor. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "ping"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"connect": TCP Connect monitor</li> <li>"http": HTTP monitor</li> <li>"ping": Ping monitor</li> <li>"program": External program monitor</li> <li>"rtsp": RTSP monitor</li> <li>"sip": SIP monitor</li> <li>"tcp_transaction": TCP transaction monitor</li> </ul> </li> </ul>
udp_accept_all	If this monitor uses UDP, should it accept responses from any IP and port? <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
use_ssl	Whether or not the monitor should connect using SSL. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
verbose	Whether or not the monitor should emit verbose logging. This is useful for diagnosing problems. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

### vtm\_monitor\_script

URI Endpoint: /api/tm/8.1/config/active/monitor\_scripts

An executable program that can be used to by external program monitors to report the health of backend services.

```
# Example usage of vtm_monitor_script resource...
resource "vtm_monitor_script" "my_monitor_script" {
  content = ""
  name = "MyMonitorScript"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_appliance\_nat

URI Endpoint: /api/tm/8.1/config/active/appliance/nat

The NAT configuration file stores rules controlling NAT on an appliance.

```
# Example usage of vtm_appliance_nat resource...
resource "vtm_appliance_nat" "my_appliance_nat" {
}
```

Property	Description
many_to_one_all_ports	<p>This is table 'many_to_one_all_ports'</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- rule_number (String): A unique rule identifier (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- pool (String): Pool of a "many to one overload" type NAT rule. (Required)</li> <li>- tip (String): TIP Group of a "many to one overload" type NAT rule. (Required)</li> </ul> </li> </ul>
many_to_one_all_ports_json	<p>A Traffic Manager REST-compatible JSON representation of the "many_to_one_all_ports" table property. Use this field with the "vtm_appliance_nat_many_to_one_all_ports_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
many_to_one_port_locked	<p>This is table 'many_to_one_port_locked'</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- rule_number (String): A unique rule identifier (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- pool (String): Pool of a "many to one port locked" type NAT rule. (Required)</li> <li>- port (UInt): Port number of a "many to one port locked" type NAT rule. (Required)</li> <li>- protocol (Enum(String)): Protocol of a "many to one port locked" type NAT rule. (Required)</li> </ul> </li> </ul> <p>Permitted values:</p> <ul style="list-style-type: none"> <li>"icmp": ICMP</li> <li>"sctp": SCTP</li> <li>"tcp": TCP</li> <li>"udp": UDP</li> <li>"udplite": UDPLITE</li> </ul> <ul style="list-style-type: none"> <li>- tip (String): TIP Group of a "many to one port locked" type NAT rule. (Required)</li> </ul>
many_to_one_port_locked_json	<p>A Traffic Manager REST-compatible JSON representation of the "many_to_one_port_locked" table property. Use this field with the "vtm_appliance_nat_many_to_one_port_locked_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
one_to_one	<p>This is table 'one_to_one'</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- rule_number (String): A unique rule identifier (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- enable_inbound (Boolean): Enabling the inbound part of a "one to one" type NAT rule. (Required)</li> <li>- ip (String): IP Address of a "one to one" type NAT rule. (Required)</li> <li>- tip (String): TIP group of a "one to one" type NAT rule. (Required)</li> </ul> </li> </ul>
one_to_one_json	<p>A Traffic Manager REST-compatible JSON representation of the "one_to_one" table property. Use this field with the "vtm_appliance_nat_one_to_one_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
port_mapping	<p>This is table 'port_mapping'</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- rule_number (String): A unique rule identifier (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- dport_first (UInt): First port of the dest. port range of a "port mapping" rule. (Required)</li> <li>- dport_last (UInt): Last port of the dest. port range of a "port mapping" rule. (Required)</li> <li>- virtual_server (String): Target Virtual Server of a "port mapping" rule. (Required)</li> </ul> </li> </ul>
port_mapping_json	<p>A Traffic Manager REST-compatible JSON representation of the "port_mapping" table property. Use this field with the "vtm_appliance_nat_port_mapping_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

## vtm\_pool

URI Endpoint: /api/tm/8.1/config/active/pools

The conf/pools directory contains configuration files for backend node pools. The name of a file is the name of the pool it defines. Pools can be configured under the Services > Pools section of the Admin Server UI or by using functions under the Pool section of the SOAP API and CLI.



```
# Example usage of vtm_pool resource...
resource "vtm_pool" "my_pool" {
  name = "MyPool"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
auto_scaling_addnode_delaytime	The time in seconds from the creation of the node which the traffic manager should wait before adding the node to the autoscaled pool. Set this to allow applications on the newly created node time to initialize before being sent traffic. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_cloud_credentials	The Cloud Credentials object containing authentication credentials to use in cloud API calls. <ul style="list-style-type: none"> <li>Type: Reference(cloud-api)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_cluster	The ESX host or ESX cluster name to put the new virtual machine instances on. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_data_center	The name of the logical datacenter on the vCenter server. Virtual machines will be scaled up and down under the datacenter root folder. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_data_store	The name of the datastore to be used by the newly created virtual machine. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_enabled	Are the nodes of this pool subject to autoscaling? If yes, nodes will be automatically added and removed from the pool by the chosen autoscaling mechanism. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
auto_scaling_external	<p>Whether or not autoscaling is being handled by an external system. Set this value to Yes if all aspects of autoscaling are handled by an external system, such as RightScale. If set to No, the traffic manager will determine when to scale the pool and will communicate with the cloud provider to create and destroy nodes as necessary.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
auto_scaling_extraargs	<p>Any extra arguments to the autoscaling API. Each argument can be separated by comma. E.g in case of EC2, it can take extra parameters to the Amazon's RunInstance API say DisableApiTermination=false,Placement.Tenancy=default.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_hysteresis	<p>The time period in seconds for which a change condition must persist before the change is actually instigated.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "20"</li> </ul>
auto_scaling_imageid	<p>The identifier for the image of the instances to create.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_ips_to_use	<p>Which type of IP addresses on the node to use. Choose private IPs if the traffic manager is in the same cloud as the nodes, otherwise choose public IPs.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "publicips"</li> <li>Permitted values: "private_ips": Private IP addresses "publicips": Public IP addresses</li> </ul>
auto_scaling_last_node_idle_time	<p>The time in seconds for which the last node in an autoscaled pool must have been idle before it is destroyed. This is only relevant if min_nodes is 0.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3600"</li> </ul>
auto_scaling_max_nodes	<p>The maximum number of nodes in this autoscaled pool.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4"</li> </ul>

Property	Description
auto_scaling_min_nodes	The minimum number of nodes in this autoscaled pool. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>
auto_scaling_name	The beginning of the name of nodes in the cloud that are part of this autoscaled pool. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_port	The port number to use for each node in this autoscaled pool. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "80"</li> </ul>
auto_scaling_refractory	The time period in seconds after the instigation of a re-size during which no further changes will be made to the pool size. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "180"</li> </ul>
auto_scaling_response_time	The expected response time of the nodes in ms. This time is used as a reference when deciding whether a node's response time is conforming. All responses from all the nodes will be compared to this reference and the percentage of conforming responses is the base for decisions about scaling the pool up or down. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1000"</li> </ul>
auto_scaling_scale_down_level	The fraction, in percent, of conforming requests above which the pool size is decreased. If the percentage of conforming requests exceeds this value, the pool is scaled down. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "95"</li> </ul>
auto_scaling_scale_up_level	The fraction, in percent, of conforming requests below which the pool size is increased. If the percentage of conforming requests drops below this value, the pool is scaled up. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "40"</li> </ul>
auto_scaling_securitygroupids	List of security group IDs to associate to the new EC2 instance. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
auto_scaling_size_id	The identifier for the size of the instances to create. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auto_scaling_subnetids	List of subnet IDs where the new EC2-VPC instance(s) will be launched. Instances will be evenly distributed among the subnets. If the list is empty, instances will be launched inside EC2-Classical. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
bandwidth_class	The Bandwidth Management Class this pool uses, if any. <ul style="list-style-type: none"> <li>Type: Reference(config-bandwidth)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_max_connect_time	How long the pool should wait for a connection to a node to be established before giving up and trying another node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4"</li> </ul>
connection_max_connections_per_node	The maximum number of concurrent connections allowed to each back-end node in this pool per machine. A value of 0 means unlimited connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_max_queue_size	The maximum number of connections that can be queued due to connections limits. A value of 0 means unlimited queue size. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_max_reply_time	How long the pool should wait for a response from the node before either discarding the request or trying another node (retryable requests only). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
connection_max_transactions_per_node	The maximum number of concurrent transactions allowed to each back-end node in this pool per machine. A value of 0 means unlimited transactions. Idle connections kept alive for reuse do not count against this limit. A transaction begins by allocating a connection for sending the request, and ends (for the purposes of queuing) after a complete response has been received from the node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
connection_queue_timeout	The maximum time to keep a connection queued in seconds. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
dns_autoscale_enabled	When enabled, the Traffic Manager will periodically resolve the hostnames in the "hostnames" list using a DNS query, and use the results to automatically add, remove or update the IP addresses of the nodes in the pool. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
dns_autoscale_hostnames	A list of hostnames which will be used for DNS-derived autoscaling <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
dns_autoscale_port	The port number to use for each node when using DNS-derived autoscaling <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "80"</li> </ul>
failure_pool	If all of the nodes in this pool have failed, then requests can be diverted to another pool. <ul style="list-style-type: none"> <li>Type: Reference(config-pool)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ftp_support_rfc_2428	Whether or not the backend IPv4 nodes understand the EPRT and EPSV command from RFC 2428. It is always assumed that IPv6 nodes support these commands. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
http_keepalive	Whether or not the pool should maintain HTTP keepalive connections to the nodes. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http_keepalive_non_idempotent	Whether or not the pool should maintain HTTP keepalive connections to the nodes for non-idempotent requests. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
kerberos_protocol_transition_principal	<p>The Kerberos principal the traffic manager should use when performing Kerberos Protocol Transition.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
kerberos_protocol_transition_target	<p>The Kerberos principal name of the service this pool targets.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
load_balancing_algorithm	<p>The load balancing algorithm that this pool uses to distribute load across its nodes.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "round_robin"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"fastest_response_time": The Response Time algorithm monitors the response times for recent requests to each node. It sends each new request to the node that has recently been responding the most quickly.</li> <li>"least_connections": This algorithm sends each new request to the node with the fewest currently active connections.</li> <li>"perceptive": The Perceptive algorithm uses a combination of response time data and connection counts to predict which node is likely to have the fastest response time for each request.</li> <li>"random": This algorithm chooses a random node for each request.</li> <li>"round_robin": This algorithm distributes traffic by assigning each request to a new node in turn.</li> <li>"weighted_least_connections": This algorithm works in a similar way to the Least Connections algorithm, but assigns more requests to nodes with a greater 'weight'.</li> <li>"weighted_round_robin": Weighted Round Robin works in a similar way to Round Robin, but assigns more requests to nodes with a greater 'weight'.</li> </ul> </li> </ul>
load_balancing_priority_enabled	<p>Enable priority lists.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
load_balancing_priority_nodes	<p>Minimum number of highest-priority active nodes.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>

Property	Description
max_connection_attempts	<p>The maximum number of nodes to which the traffic manager will attempt to send a request before returning an error to the client. Requests that are non-retryable will be attempted against only one node. Zero signifies no limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
max_idle_connections_per_node	<p>The maximum number of unused HTTP keepalive connections that should be maintained to an individual node. Zero signifies no limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "50"</li> </ul>
max_timed_out_connection_attempts	<p>The maximum number of connection attempts the traffic manager will make where the server fails to respond within the time limit defined by the max_reply_time setting. Zero signifies no limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2"</li> </ul>
monitors	<p>The monitors assigned to this pool, used to detect failures in the back end nodes.</p> <ul style="list-style-type: none"> <li>Type: Set(Reference(config-monitor))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
node_close_on_death	<p>Close all connections to a node once we detect that it has failed.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
node_close_with_rst	<p>Whether or not connections to the back-end nodes should be closed with a RST packet, rather than a FIN packet. This avoids the TIME_WAIT state, which on rare occasions allows wandering duplicate packets to be safely ignored.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
node_connection_attempts	<p>The number of times the software will attempt to connect to the same back-end node before marking it as failed. This is only used when passive_monitoring is enabled.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "3"</li> </ul>

Property	Description
node_delete_behavior	<p>Specify the deletion behavior for nodes in this pool.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "immediate"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"drain": Allow existing connections to the node to finish before deletion.</li> <li>"immediate": All connections to the node are closed immediately.</li> </ul> </li> </ul>
node_drain_to_delete_timeout	<p>The maximum time that a node will be allowed to remain in a draining state after it has been deleted. A value of 0 means no maximum time.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
node_retry_fail_time	<p>The amount of time, in seconds, that a traffic manager will wait before re-trying a node that has been marked as failed by passive monitoring.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
nodes_table	<p>A table of all nodes in this pool. A node should be specified as a &lt;ip&gt;:&lt;port&gt; pair, and has a state, weight and priority.</p> <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>node (String): A node is a combination of an ip address and port (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>priority (UInt): The priority of the node, higher values signify higher priority. If a priority is not specified for a node it is assumed to be 1.</li> <li>state (Enum(String)): The state of the pool, which can either be Active, Draining or Disabled</li> </ul> </li> </ul> <p>Permitted values:</p> <ul style="list-style-type: none"> <li>"active": The node is active.</li> <li>"disabled": The node is disabled.</li> <li>"draining": The node is draining.</li> </ul> <ul style="list-style-type: none"> <li>weight (Int): Weight for the node. The actual value in isolation does not matter: As long as it is a valid integer 1-100, the per-node weightings are calculated on the relative values between the nodes.</li> <li>source_ip (String): The source address the Traffic Manager uses to connect to this node.</li> </ul>



Property	Description
nodes_table_json	<p>A Traffic Manager REST-compatible JSON representation of the "nodes_table" table property. Use this field with the "vtm_pool_nodes_table_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
note	<p>A description of the pool.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
passive_monitoring	<p>Whether or not the software should check that 'real' requests (i.e. not those from monitors) to this pool appear to be working. This should normally be enabled, so that when a node is refusing connections, responding too slowly, or sending back invalid data, it can mark that node as failed, and stop sending requests to it. If this is disabled, you should ensure that suitable health monitors are configured to check your servers instead, otherwise failed requests will not be detected and subsequently retried.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
persistence_class	<p>The default Session Persistence class this pool uses, if any.</p> <ul style="list-style-type: none"> <li>• Type: Reference(config-persistence)</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
service_discovery_enabled	<p>Are the nodes of this pool determined by a Service Discovery plugin? If yes, nodes will be automatically added and removed from the pool by the traffic manager.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
service_discovery_interval	<p>The minimum time before rerunning the Service Discovery plugin</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "10"</li> </ul>
service_discovery_plugin	<p>The plugin script a Service Discovery autoscaled pool should use to retrieve the list of nodes.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
service_discovery_plugin_args	<p>The arguments for the script specified in "service_discovery!plugin", e.g. a common instance tag, or name of a managed group of cloud instances.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
service_discovery_timeout	<p>The maximum time a plugin should be allowed to run before timing out. Set to 0 for no timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
smtp_send_starttls	<p>If we are encrypting traffic for an SMTP connection, should we upgrade to SSL using STARTTLS.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_cipher_suites	<p>The SSL/TLS cipher suites to allow for connections to a back-end node. Leaving this empty will make the pool use the globally configured cipher suites, see configuration key <code>ssl!cipher_suites</code> in the Global Settings section of the System tab. See there for how to specify SSL/TLS cipher suites.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_client_auth	<p>Whether or not a suitable certificate and private key from the SSL Client Certificates catalog be used if the back-end server requests client authentication.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_common_name_match	<p>A list of names against which the 'common name' of the certificate is matched; these names are used in addition to the node's hostname or IP address as specified in the config file or added by the autoscaler process.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_elliptic_curves	<p>The SSL elliptic curve preference list for SSL connections from this pool using TLS version 1.0 or higher. Leaving this empty will make the pool use the globally configured preference list. The named curves P256, P384 and P521 may be configured.</p> <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ssl_enable	<p>Whether or not the pool should encrypt data before sending it to a back-end node.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_enhance	<p>SSL protocol enhancements allow your traffic manager to prefix each new SSL connection with information about the client. This enables Pulse Secure Virtual Traffic Manager virtual servers referenced by this pool to discover the original client's IP address. Only enable this if you are using nodes for this pool which are Pulse Secure vTMs, whose virtual servers have the ssl_trust_magic setting enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_middlebox_compatibility	<p>Whether or not TLS 1.3 middlebox compatibility mode as described in RFC 8446 appendix D.4 will be used in connections to pool nodes. Choosing the global setting means the value of configuration key ssl!middlebox_compatibility from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable use of middlebox compatibility</li> <li>"enabled": Enable use of middlebox compatibility</li> <li>"use_default": Use the global setting for use of middlebox compatibility</li> </ul> </li> </ul>
ssl_send_close_alerts	<p>Whether or not to send an SSL/TLS "close alert" when initiating a socket disconnection.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_server_name	<p>Whether or not the software should use the TLS 1.0 server_name extension, which may help the back-end node provide the correct certificate. Enabling this setting will force the use of at least TLS 1.0.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
ssl_session_cache_enabled	<p>Whether or not the SSL client cache will be used for this pool. Choosing the global setting means the value of the configuration key <code>ssl!client_cache!enabled</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable use of the session cache</li> <li>"enabled": Enable use of the session cache</li> <li>"use_default": Use the global setting for use of the session cache</li> </ul> </li> </ul>
ssl_session_tickets_enabled	<p>Whether or not SSL session tickets, including TLS &gt;= 1.3 PSKs, will be used for this pool if the session cache is also enabled. Choosing the global setting means the value of the configuration key <code>ssl!client_cache!tickets_enabled</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable use of session tickets</li> <li>"enabled": Enable use of session tickets</li> <li>"use_default": Use the global setting for use of session tickets</li> </ul> </li> </ul>
ssl_signature_algorithms	<p>The SSL signature algorithms preference list for SSL connections from this pool using TLS version 1.2 or higher. Leaving this empty will make the pool use the globally configured preference list, <code>signature_algorithms</code> in the <code>ssl</code> section of the <code>global_settings</code> resource. See there and in the online help for how to specify SSL signature algorithms.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_ssl_fixed_client_certificate	<p>An entry in the SSL client certificates catalog, containing a certificate and private key to be used whenever client authentication is requested. If set, this overrides server request parameters.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_strict_verify	<p>Whether or not strict certificate verification should be performed. This will turn on checks to disallow server certificates that don't match the server name or a name in the <code>ssl_common_name_match</code> list, are self-signed, expired, revoked, or have an unknown CA.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
ssl_support_ssl3	<p>Whether or not SSLv3 is enabled for this pool. Choosing the global setting means the value of the configuration key <code>ssl!support_ssl3</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable SSLv3</li> <li>"enabled": Enable SSLv3</li> <li>"use_default": Use the global setting for SSLv3</li> </ul> </li> </ul>
ssl_support_tls1	<p>Whether or not TLSv1.0 is enabled for this pool. Choosing the global setting means the value of the configuration key <code>ssl!support_tls1</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.0</li> <li>"enabled": Enable TLSv1.0</li> <li>"use_default": Use the global setting for TLSv1.0</li> </ul> </li> </ul>
ssl_support_tls1_1	<p>Whether or not TLSv1.1 is enabled for this pool. Choosing the global setting means the value of the configuration key <code>ssl!support_tls1_1</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.1</li> <li>"enabled": Enable TLSv1.1</li> <li>"use_default": Use the global setting for TLSv1.1</li> </ul> </li> </ul>
ssl_support_tls1_2	<p>Whether or not TLSv1.2 is enabled for this pool. Choosing the global setting means the value of the configuration key <code>ssl!support_tls1_2</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.2</li> <li>"enabled": Enable TLSv1.2</li> <li>"use_default": Use the global setting for TLSv1.2</li> </ul> </li> </ul>

Property	Description
ssl_support_tls1_3	<p>Whether or not TLSv1.3 is enabled for this pool. Choosing the global setting means the value of the configuration key <code>ssl!support_tls1_3</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.3</li> <li>"enabled": Enable TLSv1.3</li> <li>"use_default": Use the global setting for TLSv1.3</li> </ul> </li> </ul>
tcp_nagle	<p>Whether or not Nagle's algorithm should be used for TCP connections to the back-end nodes.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
transparent	<p>Whether or not connections to the back-ends appear to originate from the source client IP address.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
udp_accept_from	<p>The IP addresses and ports from which responses to UDP requests should be accepted. If set to accept responses from a specific set of IP addresses, you will need to enter a CIDR Mask (such as 10.100.0.0/16).</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "dest_only"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"all": Any IP address and any port.</li> <li>"dest_ip_only": Only the IP address to which the request was sent, but from any port.</li> <li>"dest_only": Only the IP address and port to which the request was sent.</li> <li>"ip_mask": Only a specific set of IP addresses, but from any port.</li> </ul> </li> </ul>
udp_accept_from_mask	<p>The CIDR mask that matches IPs we want to receive responses from.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
udp_response_timeout	<p>The maximum length of time that a node is permitted to take after receiving a UDP request packet before sending a reply packet. Zero indicates that there is no maximum, preventing a node that does not send replies from being presumed to have failed.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_protection

URI Endpoint: /api/tm/8.1/config/active/protection

A protection class specifies the level of protection against network attacks for a virtual server.

```
# Example usage of vtm_protection resource...
resource "vtm_protection" "my_protection" {
  name = "MyProtection"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
access_restriction_allowed	Always allow access to these IP addresses. This overrides the connection limits for these machines, but does not stop other restrictions such as HTTP validity checks. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
access_restriction_banned	Disallow access to these IP addresses. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
concurrent_connections_max_10_connections	Additional limit on maximum concurrent connections from the top 10 busiest connecting IP addresses combined. The value should be between 1 and 10 times the max_1_connections limit. (This limit is disabled if per_process_connection_count is No, or max_1_connections is 0, or min_connections is 0.) <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "200"</li> </ul>
concurrent_connections_max_1_connections	Maximum concurrent connections each connecting IP address is allowed. Set to 0 to disable this limit. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
concurrent_connections_min_connections	Entry threshold for the max_10_connections limit: the max_10_connections limit is not applied to connecting IP addresses with this many or fewer concurrent connections. Setting to 0 disables both the max_1_connections and max_10_connections limits, if per_process_connection_count is Yes. (If per_process_connection_count is No, this setting is ignored.) <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4"</li> </ul>

Property	Description
concurrent_connections_per_process_connection_count	<p>Whether concurrent connection counting and limits are per-process. (Each Traffic Manager typically has several processes: one process per available CPU core.) If Yes, a connecting IP address may make that many connections to each process within a Traffic Manager. If No, a connecting IP address may make that many connections to each Traffic Manager as a whole.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
connection_rate_max_connection_rate	<p>Maximum number of new connections each connecting IP address is allowed to make in the rate_timer interval. Set to 0 to disable this limit. If applied to an HTTP Virtual Server each request sent on a connection that is kept alive counts as a new connection. The rate limit is per process: each process within a Traffic Manager accepts new connections from the connecting IP address at this rate. (Each Traffic Manager typically has several processes: one process per available CPU core).</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
connection_rate_rate_timer	<p>How frequently the max_connection_rate is assessed. For example, a value of 1 (second) will impose a limit of max_connection_rate connections per second; a value of 60 will impose a limit of max_connection_rate connections per minute. The valid range is 1-99999 seconds.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "60"</li> </ul>
debug	<p>Whether or not to output verbose logging.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
enabled	<p>Enable or disable this service protection class.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
http_check_rfc2396	<p>Whether or not requests with poorly-formed URLs be should be rejected. This tests URL compliance as defined in RFC2396. Note that enabling this may block some older, non-conforming web browsers.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>



Property	Description
http_max_body_length	<p>Maximum permitted length of HTTP request body data, set to 0 to disable the limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_max_header_length	<p>Maximum permitted length of a single HTTP request header (key and value), set to 0 to disable the limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_max_request_length	<p>Maximum permitted size of all the HTTP request headers, set to 0 to disable the limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_max_url_length	<p>Maximum permitted URL length, set to 0 to disable the limit.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_reject_binary	<p>Whether or not URLs and HTTP request headers that contain binary data (after decoding) should be rejected.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
http_send_error_page	<p>This setting tells the traffic manager to send an HTTP error message if a connection fails the service protection tests, instead of just dropping it. Details of which HTTP response will be sent when particular tests fail can be found in the Help section for this page.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
log_time	<p>Log service protection messages at these intervals. If set to 0 no messages will be logged and no alerts will be sent.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
note	<p>A description of the service protection class.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
rule	<p>A TrafficScript rule that will be run on the connection after the service protection criteria have been evaluated. This rule will be executed prior to normal rules configured for the virtual server.</p> <ul style="list-style-type: none"> <li>• Type: Reference(config-trafficscript)</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
testing	<p>Place the service protection class into testing mode. (Log when this class would have dropped a connection, but allow all connections through).</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>

### vtm\_application\_firewall

URI Endpoint: /api/tm/8.1/config/active/application\_firewall

The conf/zeusafm.conf file contains configuration files for the application firewall. Some keys present in the zeusafm.conf are not documented here. Refer to the Pulse Secure Web Application Firewall documentation for further details. The configuration can be edited under the System > Application Firewall section of the Administration Server or by using functions under the AFM section of the SOAP API and CLI.

Property	Description
There are no properties to display for this resource.	

### vtm\_rate

URI Endpoint: /api/tm/8.1/config/active/rate

A rate shaping class restricts the number of connections being processed by a virtual server at once.

```
# Example usage of vtm_rate resource...
resource "vtm_rate" "my_rate" {
  name = "MyRate"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
max_rate_per_minute	Requests that are associated with this rate class will be rate-shaped to this many requests per minute, set to 0 to disable the limit. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
max_rate_per_second	Although requests will be rate-shaped to the max_rate_per_minute, the traffic manager will also rate limit per-second. This smooths traffic so that a full minute's traffic will not be serviced in the first second of the minute, set this to 0 to disable the per-second limit. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	A description of the rate class. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_rule

URI Endpoint: /api/tm/8.1/config/active/rules

TrafficScript rules allow traffic inspection and modification.

```
# Example usage of vtm_rule resource...
resource "vtm_rule" "my_rule" {
  content = ""
  name = "MyRule"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

## vtm\_service\_level\_monitor

URI Endpoint: /api/tm/8.1/config/active/service\_level\_monitors

Service level monitoring is used to produce alerts when an application's performance is degraded. This is done by monitoring the response time of connections to a virtual server.

```
# Example usage of vtm_service_level_monitor resource...
resource "vtm_service_level_monitor" "my_service_level_monitor" {
  name = "MyServiceLevelMonitor"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
note	A description for the SLM class. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
response_time	Responses that start being sent to the client within this time limit, expressed in milliseconds, are treated as conforming. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1000"</li> </ul>
serious_threshold	When the percentage of conforming responses drops below this level, a serious error level message will be emitted. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
warning_threshold	When the percentage of conforming responses drops below this level, a warning message will be emitted. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "50"</li> </ul>

## vtm\_ssl\_client\_key

URI Endpoint: /api/tm/8.1/config/active/ssl/client\_keys

SSL Client Certificates are used when connecting to backend nodes that require client certificate authentication.

```
# Example usage of vtm_ssl_client_key resource...
resource "vtm_ssl_client_key" "my_ssl_client_key" {
  name = "MySslClientKey"
  note =
  private =
  public =
  request =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
note	Notes for this certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
private	Private key for certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
public	Public certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
request	Certificate Signing Request for certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_ssl\_server\_key

URI Endpoint: /api/tm/8.1/config/active/ssl/server\_keys

SSL Server Certificates are presented to clients by virtual servers when SSL decryption is enabled.

```
# Example usage of vtm_ssl_server_key resource...
resource "vtm_ssl_server_key" "my_ssl_server_key" {
  name = "MySslServerKey"
  note =
  private =
  public =
  request =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
note	Notes for this certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
private	Private key for certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
public	Public certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
request	Certificate Signing Request for certificate <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_ssl\_ticket\_key

URI Endpoint: /api/tm/8.1/config/active/ssl/ticket\_keys

Configuration for SSL ticket encryption keys when managed externally via the ssl/ticket\_keys REST API endpoints.

```
# Example usage of vtm_ssl_ticket_key resource...
resource "vtm_ssl_ticket_key" "my_ssl_ticket_key" {
  id =
  key =
  name = "MySslTicketKey"
  validity_end =
  validity_start =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
algorithm	The algorithm used to encrypt session tickets. The algorithm determines the length of the key that must be provided. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "aes_256_cbc_hmac_sha256"</li> <li>Permitted values: "aes_256_cbc_hmac_sha256": AES-256 CBC with HMAC-SHA256. Requires a total of 64 bytes of key material.</li> </ul>
identifier	A 16-byte key identifier, with each byte encoded as two hexadecimal digits. Key identifiers are transmitted in plaintext at the beginning of a TLS session ticket, and are used to identify the ticket encryption key that was used to encrypt a ticket. (They correspond to the 'key_name' field in RFC 5077.) They are required to be unique across the set of SSL ticket encryption keys. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
key	The session ticket encryption key, with each byte encoded as two hexadecimal digits. The required key length is determined by the chosen key algorithm. See the documentation for the 'algorithm' field for more details. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
validity_end	The latest time at which this key may be used to encrypt new session tickets. Given as number of seconds since the epoch (1970-01-01T00:00:00Z). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
validity_start	The earliest time at which this key may be used to encrypt new session tickets. Given as number of seconds since the epoch (1970-01-01T00:00:00Z). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_ssl\_ca

URI Endpoint: /api/tm/8.1/config/active/ssl/cas

SSL certificate authority certificates (CAs) and certificate revocation lists (CRLs) can be used when validating server and client certificates.

```
# Example usage of vtm_ssl_ca resource...
resource "vtm_ssl_ca" "my_ssl_ca" {
  content = ""
  name = "MySslCa"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_security

URI Endpoint: /api/tm/8.1/config/active/security

Security settings that restrict remote administration for the cluster. Additional security options can be found in Global Settings.

```
# Example usage of vtm_security resource...
resource "vtm_security" "my_security" {
}
```

Property	Description
access	Access to the admin server and REST API is restricted by usernames and passwords. You can further restrict access to just trusted IP addresses, CIDR IP subnets or DNS wildcards. These access restrictions are also used when another traffic manager initially joins the cluster, after joining the cluster these restrictions are no longer used. Care must be taken when changing this setting, as it can cause the administration server to become inaccessible. Access to the admin UI will not be affected until it is restarted. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssh_intrusion_bantime	The amount of time in seconds to ban an offending host for. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "600"</li> </ul>



Property	Description
ssh_intrusion_blacklist	The list of hosts to permanently ban, identified by IP address or DNS hostname in a space-separated list. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssh_intrusion_enabled	Whether or not the SSH Intrusion Prevention tool is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssh_intrusion_findtime	The window of time in seconds the maximum number of connection attempts applies to. More than (maxretry) failed attempts in this time span will trigger a ban. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "600"</li> </ul>
ssh_intrusion_maxretry	The number of failed connection attempts a host can make before being banned. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "6"</li> </ul>
ssh_intrusion_whitelist	The list of hosts to never ban, identified by IP address, DNS hostname or subnet mask, in a space-separated list. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_servicediscovery

URI Endpoint: /api/tm/8.1/config/active/servicediscovery

The conf/servicediscovery directory contains plugins for use with Service Discovery for pool nodes.

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
content	Object content <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>

### vtm\_persistence

URI Endpoint: /api/tm/8.1/config/active/persistence

A session persistence class is used to identify the session a new connection belongs too and deliver it to the same backend node.

```
# Example usage of vtm_persistence resource...
resource "vtm_persistence" "my_persistence" {
  name = "MyPersistence"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
cookie	The cookie name to use for tracking session persistence. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
delete	Whether or not the session should be deleted when a session failure occurs. (Note, setting a failure mode of 'choose a new node' implicitly deletes the session.) <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
failure_mode	The action the pool should take if the session data is invalid or it cannot contact the node specified by the session. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "new_node"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"close": Close the connection (using error_file on Virtual Servers &gt; Edit &gt; Protocol Settings)</li> <li>"new_node": Choose a new node to use</li> <li>"url": Redirect the user to a given URL</li> </ul> </li> </ul>
note	A description of the session persistence class. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
subnet_prefix_length_v4	When using IP-based session persistence, ensure all requests from this IPv4 subnet, specified as a prefix length, are sent to the same node. If set to 0, requests from different IPv4 addresses will be load-balanced individually. <ul style="list-style-type: none"> <li>Type: Int</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
subnet_prefix_length_v6	<p>When using IP-based session persistence, ensure all requests from this IPv6 subnet, specified as a prefix length, are sent to the same node. If set to 0, requests from different IPv6 addresses will be load-balanced individually.</p> <ul style="list-style-type: none"> <li>Type: Int</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
transparent_always_set_cookie	<p>Whether or not the cookie should be inserted in every response sent to the client when using transparent session affinity. If set to No then the cookie is inserted only if the corresponding request did not already contain a matching cookie.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
transparent_directives	<p>The cookie directives to include in the cookie sent when using transparent session affinity. If more than one directive is included, the semi-colon separator between them must be included in this string. The semi-colon separator between the cookie value and the first directive should not be included in this string.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
type	<p>The type of session persistence to use.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "ip"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"asp": ASP and ASP.NET session persistence</li> <li>"cookie": Monitor application cookies</li> <li>"ip": IP-based persistence</li> <li>"j2ee": J2EE session persistence</li> <li>"named": Named Node session persistence</li> <li>"ssl": SSL Session ID persistence</li> <li>"transparent": Transparent session affinity</li> <li>"universal": Universal session persistence</li> <li>"x_zeus": X-Zeus-Backend cookies</li> </ul> </li> </ul>
url	<p>The redirect URL to send clients to if the session persistence is configured to redirect users when a node dies.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

### vtm\_traffic\_ip\_group

URI Endpoint: /api/tm/8.1/config/active/traffic\_ip\_groups

Traffic IP groups are sets of IP addresses that are distributed across a cluster for fault tolerance.

```
# Example usage of vtm_traffic_ip_group resource...
resource "vtm_traffic_ip_group" "my_traffic_ip_group" {
  name = "MyTrafficIpGroup"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
enabled	If set to No, the traffic IP group will be disabled and none of the traffic IP addresses will be raised. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
hash_source_port	Whether or not the source port should be taken into account when deciding which traffic manager should handle a request. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ip_assignment_mode	Configure how traffic IPs are assigned to traffic managers in Single-Hosted mode <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "balanced"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"alphabetic": Alphabetical order of traffic manager hostnames</li> <li>"balanced": Approximately balanced between traffic managers</li> </ul> </li> </ul>
ip_mapping	A table assigning traffic IP addresses to machines that should host them. Traffic IP addresses not specified in this table will automatically be assigned to a machine. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>ip (String): A traffic IP address (from the ipaddresses property). (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>traffic_manager (String): The name of the traffic manager that should host the IP address. (Required)</li> </ul> </li> </ul>
ip_mapping_json	A Traffic Manager REST-compatible JSON representation of the "ip_mapping" table property. Use this field with the "vtm_traffic_ip_group_ip_mapping_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ipaddresses	The IP addresses that belong to the Traffic IP group. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
keptogether	If set to Yes then all the traffic IPs will be raised on a single traffic manager. By default they're distributed across all active traffic managers in the traffic IP group. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
location	The location in which the Traffic IP group is based. <ul style="list-style-type: none"> <li>Type: Int</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
machines	The traffic managers that can host the traffic IP group's IP addresses. <ul style="list-style-type: none"> <li>Type: Set(Reference(config-tm))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
mode	The method used to distribute traffic IPs across machines in the cluster. If "multihosted" is used then multicast must be set to an appropriate multicast IP address. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "singlehosted"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"ec2elastic": Use an EC2-Classical Elastic IP address.</li> <li>"ec2vpcelastic": Use an EC2-VPC Elastic IP address.</li> <li>"ec2vpcprivate": Use an EC2-VPC Private IP address.</li> <li>"gceexternal": Use GCE External IP addresses.</li> <li>"multihosted": Raise each address on every machine in the group (Multi-Hosted mode) - IPv4 only</li> <li>"rhi": Use route health injection to route traffic to the active machine - IPv4 only</li> <li>"singlehosted": Raise each address on a single machine (Single-Hosted mode)</li> </ul> </li> </ul>
multicast	The multicast IP address used to duplicate traffic to all traffic managers in the group. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
note	A note, used to describe this Traffic IP Group <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rhi_bgp_metric_base	The base BGP routing metric for this Traffic IP group. This is the advertised routing cost for the active traffic manager in the cluster. It can be used to set up inter-cluster failover. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
rhi_bgp_passive_metric_offset	The BGP routing metric offset for this Traffic IP group. This is the difference between the advertised routing cost for the active and passive traffic manager in the cluster. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
rhi_ospfv2_metric_base	The base OSPFv2 routing metric for this Traffic IP group. This is the advertised routing cost for the active traffic manager in the cluster. It can be used to set up inter-cluster failover. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
rhi_ospfv2_passive_metric_offset	The OSPFv2 routing metric offset for this Traffic IP group. This is the difference between the advertised routing cost for the active and passive traffic manager in the cluster. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
rhi_protocols	A list of protocols to be used for RHI. Currently must be 'ospf' or 'bgp' or both. The default, if empty, is 'ospf', which means that it is not possible to specify no protocol. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "ospf"</li> </ul>
slaves	A list of traffic managers that are in 'passive' mode. This means that in a fully working environment, they will not have any traffic IP addresses assigned to them. <ul style="list-style-type: none"> <li>Type: Set(Reference(config-tm))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_traffic\_manager

URI Endpoint: /api/tm/8.1/config/active/traffic\_managers

The `conf/zxtms` directory contains a configuration file for each traffic manager in your cluster. The name of each file is the hostname of the traffic manager it represents. These files contain host-specific configuration data and on each installation of the software, the `conf/./global.cfg` file is sym-linked to the host's own configuration in the `conf/zxtms` directory. The files may contain a variety of configuration options that are configured in various locations under the System section of the Admin Server UI and the System section of the SOAP API and CLI.

```
# Example usage of vtm_traffic_manager resource...
resource "vtm_traffic_manager" "my_traffic_manager" {
  name = "MyTrafficManager"
}
```

Property	Description
<code>name</code>	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
<code>adminmasterxmlip</code>	The Application Firewall master XML IP. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "0.0.0.0"</li> </ul>
<code>adminslavexmlip</code>	The Application Firewall slave XML IP. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "0.0.0.0"</li> </ul>
<code>appliance_card</code>	The table of network cards of a hardware appliance <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): Network card PCI ID (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>interfaces (List(String)): The order of the interfaces of a network card (Required)</li> <li>label (String): The labels of the installed network cards</li> </ul> </li> </ul>
<code>appliance_card_json</code>	A Traffic Manager REST-compatible JSON representation of the "appliance_card" table property. Use this field with the "vtm_traffic_manager_appliance_card_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
appliance_disable_kpti	<p>Whether the traffic manager appliance should run without kernel page table isolation (KPTI). KPTI provides protection to prevent unprivileged software from being potentially able to read arbitrary memory from the kernel (i.e. the Meltdown attack, CVE-2017-5754); however this protection incurs a general system performance penalty. If you are running trusted software on the appliance, and the trade-off between performance at the cost of 'defense in depth' favors the former in your deployment, you may wish to enable this configuration key. If you are unsure, it is recommended that you leave this key disabled, which is also the default.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
appliance_dnscache	<p>The DNS cache setting the appliance should use and place in /etc/systemd/resolved.conf.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
appliance_dnssec	<p>The DNSSEC setting the appliance should use and place in /etc/systemd/resolved.conf.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "no"</li> <li>• Permitted values:  "allow_downgrade": Use DNSSEC when available  "no": DNSSEC disabled  "yes": DNSSEC enabled</li> </ul>
appliance_gateway_ipv4	<p>The default gateway.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
appliance_gateway_ipv6	<p>The default IPv6 gateway.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
appliance_hostname	<p>Name (hostname.domainname) of the appliance.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>



Property	Description
appliance_hosts	<p>A table of hostname to static ip address mappings, to be placed in the /etc/hosts file.</p> <ul style="list-style-type: none"><li>• Type: Table</li><li>• Required: false</li><li>• Primary key:<ul style="list-style-type: none"><li>- name (String): The name of a host. (Required)</li></ul></li><li>• Sub keys:<ul style="list-style-type: none"><li>- ip_address (String): The static IP address of the host. (Required)</li></ul></li></ul>
appliance_hosts_json	<p>A Traffic Manager REST-compatible JSON representation of the "appliance_hosts" table property. Use this field with the "vtm_traffic_manager_hosts_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"><li>• Type: String</li><li>• Required: false</li><li>• Default value: &lt;none&gt;</li></ul>

Property	Description
appliance_if	<p>A table of network interface specific settings.</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- name (String): A network interface name. (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- autoneg (Boolean): Whether auto-negotiation should be enabled for the interface.</li> <li>- bmode (Enum(String)): The trunking mode used for the interface (only 802.3ad is currently supported). Permitted values: <ul style="list-style-type: none"> <li>"802_3ad": IEEE 802.3ad</li> <li>"balance_alb": Adaptive Load Balancing</li> </ul> </li> <li>- bond (String): The trunk of which the interface should be a member.</li> <li>- duplex (Boolean): Whether full-duplex should be enabled for the interface.</li> <li>- mode (Enum(String)): Set the configuration mode of an interface, the interface name is used in place of the * (asterisk). Permitted values: <ul style="list-style-type: none"> <li>"dhcp": DHCP</li> <li>"static": Static</li> </ul> </li> <li>- mtu (UInt): The maximum transmission unit (MTU) of the interface.</li> <li>- speed (Enum(String)): The speed of the interface. Permitted values: <ul style="list-style-type: none"> <li>"10": 10Mbps</li> <li>"100": 100Mbps</li> <li>"1000": 1Gbps</li> <li>"10000": 10Gbps</li> <li>"100000": 100Gbps</li> <li>"40000": 40Gbps</li> </ul> </li> </ul> </li> </ul>
appliance_if_json	<p>A Traffic Manager REST-compatible JSON representation of the "appliance_if" table property. Use this field with the "vtm_traffic_manager_if_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
appliance_ip	<p>A table of network interfaces and their network settings.</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- name (String): A network interface name. (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- addr (String): The IP address for the interface. (Required)</li> <li>- isexternal (Boolean): Whether the interface is externally facing.</li> <li>- mask (String): The IP mask (netmask) for the interface. (Required)</li> </ul> </li> </ul>
appliance_ip_json	<p>A Traffic Manager REST-compatible JSON representation of the "appliance_ip" table property. Use this field with the "vtm_traffic_manager_ip_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
appliance_ipmi_lan_access	<p>Whether IPMI LAN access should be enabled or not.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>
appliance_ipmi_lan_addr	<p>The IP address of the appliance IPMI LAN channel.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
appliance_ipmi_lan_gateway	<p>The default gateway of the IPMI LAN channel.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: "0.0.0.0"</li> </ul>
appliance_ipmi_lan_ipsrc	<p>The addressing mode the IPMI LAN channel operates.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "static"</li> <li>• Permitted values: <ul style="list-style-type: none"> <li>"dhcp": Address obtained by DHCP</li> <li>"static": Static IP Address</li> </ul> </li> </ul>
appliance_ipmi_lan_mask	<p>Set the IP netmask for the IPMI LAN channel.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
appliance_ipv4_forwarding	<p>Whether or not IPv4 forwarding is enabled.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>

Property	Description
appliance_ipv6_forwarding	Whether or not IPv6 forwarding is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
appliance_licence_agreed	Whether or not the license agreement has been accepted. This determines whether or not the Initial Configuration wizard is displayed. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
appliance_manageazureroutes	Whether or not the software manages the Azure policy routing. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_manageec2conf	Whether or not the software manages the EC2 config. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_managegceroutes	Whether or not the software manages the GCE routing. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_manageiptrans	Whether or not the software manages the IP transparency <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_managereservedports	Whether or not the software manages the system configuration for reserved ports <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_managereturnpath	Whether or not the software manages return path routing. If disabled, the appliance won't modify iptables / rules / routes for this feature. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_manageservices	Whether or not the software manages the system services <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_managevpcconf	Whether or not the software manages the EC2-VPC secondary IPs. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
appliance_name_servers	The IP addresses of the nameservers the appliance should use and place in /etc/systemd/resolved.conf. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
appliance_ntpservers	The NTP servers the appliance should use to synchronize its clock. <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: "0.zeus.pool.ntp.org 1.zeus.pool.ntp.org 2.zeus.pool.ntp.org 3.zeus.pool.ntp.org"</li> </ul>
appliance_routes	A table of destination IP addresses and routing details to reach them. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): A destination IP address. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>gw (String): The gateway IP to configure for the route. (Required)</li> <li>if (String): The network interface to configure for the route. (Required)</li> <li>mask (String): The netmask to apply to the IP address. (Required)</li> </ul> </li> </ul>
appliance_routes_json	A Traffic Manager REST-compatible JSON representation of the "appliance_routes" table property. Use this field with the "vtm_traffic_manager_routes_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
appliance_search_domains	The search domains the appliance should use and place in /etc/systemd/resolved.conf. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
appliance_ssh_enabled	Whether or not the SSH server is enabled on the appliance. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
appliance_ssh_password_allowed	Whether or not the SSH server allows password based login. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
appliance_ssh_port	The port that the SSH server should listen on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "22"</li> </ul>
appliance_sysctl	Custom kernel parameters applied by the user with sysctl interface <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>sysctl (String): The name of the kernel parameter, e.g. net.ipv4.forward (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>description (String): Associated optional description for the sysctl</li> <li>value (String): The value of the kernel parameter (Required)</li> </ul> </li> </ul>
appliance_sysctl_json	A Traffic Manager REST-compatible JSON representation of the "appliance_sysctl" table property. Use this field with the "vtm_traffic_manager_appliance_sysctl_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
appliance_timezone	The timezone the appliance should use. This must be a path to a timezone file that exists under /usr/share/zoneinfo/. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "US/Pacific"</li> </ul>
appliance_vlans	The VLANs the software should raise. A VLAN should be configured using the format <dev>.<vlanid>, where <dev> is the name of a network device that exists in the host system, eth0.100 for example. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
authenticationserverip	The Application Firewall Authentication Server IP. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "0.0.0.0"</li> </ul>
cloud_platform	Cloud platform where the traffic manager is running. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
cluster_comms_allow_update	<p>Whether or not this instance of the software can send configuration updates to other members of the cluster. When not clustered this key is ignored. When clustered the value can only be changed by another machine in the cluster that has allow_update set to true. If set to false then it will not be possible to log into the admin server for this instance.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
cluster_comms_bind_ip	<p>The IP address that the software should bind to for internal administration communications. See also port. If the software is not part of a cluster the default is to use 127.0.0.1 and there should be no reason to touch this setting. If the software is part of a cluster then the default is to listen on all raised IPs, in this case an alternative configuration is to listen on a single IP address. This may be useful if you have a separate management network and wish to restrict control messages to it. It is important to ensure that the allowed_update_hosts (in the Global Settings resource) is compatible with the IP configured here.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "*"</li> </ul>
cluster_comms_external_ip	<p>This is the optional external ip of the traffic manager, which is used to circumvent natting when traffic managers in a cluster span different networks.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cluster_comms_port	<p>The port that the software should listen on for internal administration communications. See also bind_ip.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "9080"</li> </ul>
ec2_trafficips_public_enis	<p>List of MAC addresses of interfaces which the traffic manager can use to associate the EC2 elastic IPs (Traffic IPs) to the instance.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
fault_tolerance_bgp_router_id	<p>The BGP router id If set to empty, then the IPv4 address used to communicate with the default IPv4 gateway is used instead. Specifying 0.0.0.0 will stop the traffic manager routing software from running the BGP protocol.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
fault_tolerance_ospfv2_ip	<p>The traffic manager's permanent IPv4 address which the routing software will use for peering and transit traffic, and as its OSPF router ID. If set to empty, then the address used to communicate with the default IPv4 gateway is used instead. Specifying 0.0.0.0 will stop the traffic manager routing software from running the OSPF protocol.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
fault_tolerance_ospfv2_neighbors_addrs	<p>The IP addresses of routers which are expected to be found as OSPFv2 neighbors of the traffic manager. A warning will be reported if some of the expected routers are not peered, and an error will be reported if none of the expected routers are peered. An empty list disables monitoring. The special value %gateway% is a placeholder for the default gateway.</p> <ul style="list-style-type: none"> <li>• Type: Set(String)</li> <li>• Required: false</li> <li>• Default value: "%gateway%"</li> </ul>
iptables_config_enabled	<p>Whether the Traffic Manager should configure the iptables built-in chains to call Traffic Manager defined rules (e.g. the IP transparency chain). This should only be disabled in case of conflict with other software that manages iptables, e.g. firewalls. When disabled, you will need to add rules manually to use these features - see the user manual for details.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
iptrans_fwmark	<p>The netfilter forwarding mark to use for IP transparency rules</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "320"</li> </ul>
iptrans_iptables_enabled	<p>Whether IP transparency may be used via netfilter/iptables. This requires Linux 2.6.24 and the iptables socket extension. For older Linux versions, the "ztrans" kernel module may be used instead.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: true</li> </ul>
iptrans_routing_table	<p>The special routing table ID to use for IP transparency rules</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "320"</li> </ul>
java_port	<p>The port the Java Extension handler process should listen on. This port will be bound for localhost communications only.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "9060"</li> </ul>



Property	Description
location	This is the location of the local traffic manager is in. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
nameip	Replace Traffic Manager name with an IP address. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
num_optimizer_threads	How many worker threads the Web Accelerator process should create to optimise content. By default, one thread will be created for each CPU on the system. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
num_children	The number of worker processes the software will run. By default, one child process will be created for each CPU on the system. You may wish to reduce this to effectively "reserve" CPU(s) for other processes running on the host system. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
numberofcpus	The number of Application Firewall decider process to run. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_email_address	The e-mail address sent as part of a remote licensing request. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
remote_licensing_message	A free-text field sent as part of a remote licensing request. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rest_api_bind_ips	A list of IP Addresses which the REST API will listen on for connections. The list should contain IP addresses (IPv4 or IPv6) or a single entry containing an asterisk (*). This indicates that the REST API should listen on all IP Addresses. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: "*" </li></ul>
rest_api_port	The port on which the REST API should listen for requests. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "9070"</li> </ul>

Property	Description
restserverport	The Application Firewall REST Internal API port, this port should not be accessed directly <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
snmp_allow	Restrict which IP addresses can access the SNMP command responder service. The value can be all, localhost, or a list of IP CIDR subnet masks. For example 10.100.0.0/16 would allow connections from any IP address beginning with 10.100. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: "all"</li> </ul>
snmp_auth_password	The authentication password. Required (minimum length 8 characters) if security_level includes authentication. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
snmp_bind_ip	The IP address the SNMP service should bind its listen port to. The value * (asterisk) means SNMP will listen on all IP addresses. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "*"</li> </ul>
snmp_community	The community string required for SNMPv1 and SNMPv2c commands. (If empty, all SNMPv1 and SNMPv2c commands will be rejected). <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "public"</li> </ul>
snmp_enabled	Whether or not the SNMP command responder service should be enabled on this traffic manager. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
snmp_hash_algorithm	The hash algorithm for authenticated SNMPv3 communications. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "md5"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"md5": MD5</li> <li>"sha1": SHA-1</li> </ul> </li> </ul>

Property	Description
snmp_port	<p>The port the SNMP command responder service should listen on. The value default denotes port 161 if the software is running with root privileges, and 1161 otherwise.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "default"</li> </ul>
snmp_priv_password	<p>The privacy password. Required (minimum length 8 characters) if security_level includes privacy (message encryption).</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
snmp_security_level	<p>The security level for SNMPv3 communications.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "noauthnopriv"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"authnopriv": Authentication only</li> <li>"authpriv": Authentication and Privacy</li> <li>"noauthnopriv": No Authentication, No Privacy</li> </ul> </li> </ul>
snmp_username	<p>The username required for SNMPv3 commands. (If empty, all SNMPv3 commands will be rejected).</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
trafficip	<p>A table mapping interfaces to networks, used by the traffic manager to select which interface to raise a Traffic IP on.</p> <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): A network interface. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>networks (Set(String)): A set of IP/masks to which the network interface maps. (Required)</li> </ul> </li> </ul>
trafficip_json	<p>A Traffic Manager REST-compatible JSON representation of the "trafficip" table property. Use this field with the "vtm_traffic_manager_trafficip_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
updaterip	<p>The Application Firewall Updater IP.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "0.0.0.0"</li> </ul>

## vtm\_rule\_authenticator

URI Endpoint: /api/tm/8.1/config/active/rule\_authenticators

TrafficScript authenticators define remote authentication services that can be queried via a TrafficScript rule.

```
# Example usage of vtm_rule_authenticator resource...
resource "vtm_rule_authenticator" "my_rule_authenticator" {
  name = "MyRuleAuthenticator"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
host	The hostname or IP address of the remote authenticator. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_attributes	A list of attributes to return from the search. If blank, no attributes will be returned. If set to '*' then all user attributes will be returned. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_bind_dn	The distinguished name (DN) of the 'bind' user. The traffic manager will connect to the LDAP server as this user when searching for user records. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_bind_password	The password for the bind user. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_filter	The filter used to locate the LDAP record for the user being authenticated. Any occurrences of '%u' in the filter will be replaced by the name of the user being authenticated. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_filter_base_dn	The base distinguished name (DN) under which user records are located on the server. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ldap_ssl_cert	The SSL certificate that the traffic manager should use to validate the remote server. If no certificate is specified then no signature validation will be performed. <ul style="list-style-type: none"> <li>Type: Reference(config-ssl-cacrl)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_ssl_enabled	Whether or not to enable SSL encryption to the LDAP server. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ldap_ssl_type	The type of LDAP SSL encryption to use. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "ldaps"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"ldaps": LDAPS</li> <li>"starttls": Start TLS</li> </ul> </li> </ul>
note	A description of the authenticator. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
port	The port on which the remote authenticator should be contacted. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "389"</li> </ul>

## vtm\_saml\_trustedidp

URI Endpoint: /api/tm/8.1/config/active/saml/trustedidps

Configuration for SAML IDP trust relationships.

```
# Example usage of vtm_saml_trustedidp resource...
resource "vtm_saml_trustedidp" "my_saml_trustedidp" {
  certificate =
  entity_id =
  name = "MySamlTrustedidp"
  url =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
add_zlib_header	Whether or not to add the zlib header when compressing the AuthnRequest <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
certificate	The certificate used to verify Assertions signed by the identity provider <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
entity_id	The entity id of the IDP <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
strict_verify	Whether or not SAML responses will be verified strictly <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
url	The IDP URL to which Authentication Requests should be sent <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>

## vtm\_user\_authenticator

URI Endpoint: /api/tm/8.1/config/active/user\_authenticators

A user authenticator is used to allow access to the UI and REST API by querying a remote authentication service.

```
# Example usage of vtm_user_authenticator resource...
resource "vtm_user_authenticator" "my_user_authenticator" {
  name = "MyUserAuthenticator"
  type =
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
description	A description of the authenticator. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
enabled	Whether or not this authenticator is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ldap_base_dn	The base DN (Distinguished Name) under which directory searches will be applied. The entries for your users should all appear under this DN. An example of a typical base DN is: OU=users, DC=mycompany, DC=local <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_bind_dn	Template to construct the bind DN (Distinguished Name) from the username. The string %u will be replaced by the username. Examples: %u@mycompany.local for Active Directory or cn=%u, dc=mycompany, dc=local for both LDAP and Active Directory. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_dn_method	The bind DN (Distinguished Name) for a user can either be searched for in the directory using the base distinguished name and filter values, or it can be constructed from the username. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"construct": Construct</li> <li>"none": No setting configured</li> <li>"search": Search</li> </ul> </li> </ul>
ldap_fallback_group	If the group attribute is not defined, or returns no results for the user logging in, the group named here will be used. If not specified, users will be denied access to the traffic manager if no groups matching a Permission Group can be found for them in the directory. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ldap_filter	<p>A filter that can be used to extract a unique user record located under the base DN (Distinguished Name). The string %u will be replaced by the username. This filter is used to find a user's bind DN when dn_method is set to "Search", and to extract group information if the group filter is not specified. Examples: sAMAccountName=%u for Active Directory, or uid=%u for some Unix LDAP schemas.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ldap_group_attribute	<p>The LDAP attribute that gives a user's group. If there are multiple entries for the attribute all will be extracted and they'll be lexicographically sorted, then the first one to match a Permission Group name will be used.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ldap_group_field	<p>The sub-field of the group attribute that gives a user's group. For example, if group_attribute is memberOf and this retrieves values of the form CN=mygroup, OU=groups, OU=users, DC=mycompany, DC=local you would set group_field to CN. If there are multiple matching fields only the first matching field will be used.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ldap_group_filter	<p>If the user record returned by filter does not contain the required group information you may specify an alternative group search filter here. This will usually be required if you have Unix/POSIX-style user records. If multiple records are returned the list of group names will be extracted from all of them. The string %u will be replaced by the username. Example: (&amp;(memberUid=%u)(objectClass=posixGroup))</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ldap_port	<p>The port to connect to the LDAP server on.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: "389"</li> </ul>
ldap_search_dn	<p>The bind DN (Distinguished Name) to use when searching the directory for a user's bind DN. You can leave this blank if it is possible to perform the bind DN search using an anonymous bind.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>



Property	Description
ldap_search_password	<p>If binding to the LDAP server using search_dn requires a password, enter it here.</p> <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_server	<p>The IP or hostname of the LDAP server.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ldap_ssl	<p>The type of TLS encryption, if any, to use. Usually STARTTLS will be used with port 389, and LDAPS with port 636. A Certificate Authority that the LDAP server's certificate chains back to must be present in the "Admin Certificate Authorities and Certificate Revocation Lists Catalog" under "SSL catalogs", otherwise the connection will fail.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: "ldaps": LDAPS "none": None "starttls": STARTTLS</li> </ul>
ldap_timeout	<p>Connection timeout in seconds.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
radius_fallback_group	<p>If no group is found using the vendor and group identifiers, or the group found is not valid, the group specified here will be used.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
radius_group_attribute	<p>The RADIUS identifier for the attribute that specifies an account's group. May be left blank if fallback group is specified.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>
radius_group_vendor	<p>The RADIUS identifier for the vendor of the RADIUS attribute that specifies an account's group. Leave blank if using a standard attribute (i.e. for Filter-Id set group_attribute to 11).</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "7146"</li> </ul>

Property	Description
radius_nas_identifier	This value is sent to the RADIUS server. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
radius_nas_ip_address	This value is sent to the RADIUS server, if left blank the address of the interfaced used to connect to the server will be used. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
radius_port	The port to connect to the RADIUS server on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1812"</li> </ul>
radius_secret	Secret key shared with the RADIUS server. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
radius_server	The IP or hostname of the RADIUS server. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
radius_timeout	Connection timeout in seconds. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
tacacs_plus_auth_type	Authentication type to use. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "pap"</li> <li>Permitted values: "ascii": ASCII "pap": PAP</li> </ul>
tacacs_plus_fallback_group	If group_service is not used, or no group value is provided for the user by the TACACS+ server, the group specified here will be used. If this is not specified, users with no TACACS+ defined group will be denied access. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
tacacs_plus_group_field	The TACACS+ "service" field that provides each user's group. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "permission-group"</li> </ul>

Property	Description
tacacs_plus_group_service	The TACACS+ "service" that provides each user's group field. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "zeus"</li> </ul>
tacacs_plus_port	The port to connect to the TACACS+ server on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "49"</li> </ul>
tacacs_plus_secret	Secret key shared with the TACACS+ server. <ul style="list-style-type: none"> <li>Type: Password</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
tacacs_plus_server	The IP or hostname of the TACACS+ server. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
tacacs_plus_timeout	Connection timeout in seconds. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
type	The type and protocol used by this authentication service. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"ldap": LDAP</li> <li>"radius": RADIUS</li> <li>"tacacs_plus": TACACS+</li> </ul> </li> </ul>

## vtm\_user\_group

URI Endpoint: /api/tm/8.1/config/active/user\_groups

Permission groups specify permissions for groups of users. These groups can be given read-write or read-only access to different parts of the configuration hierarchy. Each group will contain a table of permissions. Each table entry has a name that corresponds to a part of the configuration hierarchy, and a corresponding access level. The access level may have values of either none, ro (read only, this is the default), or full. Some permissions have sub-permissions, these are denoted by following the parent permission name with a colon (:) followed by the sub-permission name. The built-in admin group has a special permission key of all with the value full, this must not be altered for the admin group but can be used in other group configuration files to change the default permission level for the group.

```
# Example usage of vtm_user_group resource...
resource "vtm_user_group" "my_user_group" {
  name = "MyUserGroup"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
description	A description for the group. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
password_expire_time	Members of this group must renew their passwords after this number of days. To disable password expiry for the group set this to 0 (zero). Note that this setting applies only to local users. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
permissions	A table defining which level of permission this group has for specific configuration elements. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): Configuration element to which this group has a level of permission. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>access_level (String): Permission level for the configuration element (none, ro or full) (Required)</li> </ul> </li> </ul>
permissions_json	A Traffic Manager REST-compatible JSON representation of the "permissions" table property. Use this field with the "vtm_user_group_permissions_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
timeout	Inactive UI sessions will timeout after this number of seconds. To disable inactivity timeouts for the group set this to 0 (zero). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>

## vtm\_virtual\_server

URI Endpoint: /api/tm/8.1/config/active/virtual\_servers

The `conf/vservers` directory contains configuration files that define virtual servers. The name of a file is the name of the virtual server it defines. Virtual servers can be configured under the Services > Virtual Servers section of the Admin Server UI or by using functions under the `VirtualServer` section of the SOAP API and CLI.

```
# Example usage of vtm_virtual_server resource...
resource "vtm_virtual_server" "my_virtual_server" {
  name = "MyVirtualServer"
  pool =
  port =
}
```

Property	Description
<code>name</code>	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
<code>optimizer_enabled</code>	Whether the virtual server should optimize web content. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
<code>optimizer_profile</code>	A table of Optimizer profiles and the application scopes that apply to them. <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li><code>name</code> (String): The name of an Optimizer acceleration profile. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li><code>urls</code> (Set(String)): The application scopes which apply to the acceleration profile. (Required)</li> </ul> </li> </ul>
<code>optimizer_profile_json</code>	A Traffic Manager REST-compatible JSON representation of the "optimizer_profile" table property. Use this field with the "vtm_virtual_server_profile_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
<code>auth_saml_idp</code>	Name of the Trusted Identity Provider configuration to use. To create Identity Providers, please visit section Trusted Identity Providers <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
auth_saml_nameid_format	<p>The NameID format to request and expect from the identity provider.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"emailaddress": emailAddress</li> <li>"none": none</li> <li>"unspecified": unspecified</li> </ul> </li> </ul>
auth_saml_sp_acs_url	<p>The 'Assertion Consumer Service' endpoint for the SAML service provider on this virtual server, ie the endpoint to which the identity provider will cause the user agent to send SAML assertions. This should be an HTTPS URL, must be in the same cookie domain as all hostnames used by the end user to access the virtual server (see cookie configuration) and the port must be the port on which this virtual server is listening. It must match the URI placed by the identity provider in the 'Recipient' attribute in the SAML assertion, if present.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auth_saml_sp_entity_id	<p>The entity ID to be used by the SAML service provider function on this virtual server. This should usually be a URL, or a URN, however it may be any string. It must match the entity ID placed by the identity provider in the 'Audience' field in the SAML assertion.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
auth_saml_time_tolerance	<p>Time tolerance on authentication checks. When checking time-stamps and expiry dates against the current time on the system, allow a tolerance of this many seconds. For example, if a SAML response contains a 'NotOnOrAfter' that is 4 seconds in the past according to the local time, and the tolerance is set to 5 seconds, it will still be accepted. This is to prevent a lack of clock synchronization from resulting in rejection of SAML responses.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "5"</li> </ul>
auth_session_cookie_attributes	<p>Attributes of cookie used for authentication session.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "HttpOnly; SameSite=Strict"</li> </ul>
auth_session_cookie_name	<p>Name of cookie used for authentication session.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "VS_SamISP_Auth"</li> </ul>

Property	Description
auth_session_log_external_state	Whether or not to include state of authentication sessions stored encrypted on the client as plaintext in the logs. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
auth_session_timeout	Timeout on authentication session. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "7200"</li> </ul>
auth_type	Type of authentication to apply to requests to the virtual server. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"none": None</li> <li>"saml_sp": SAML Service Provider</li> </ul> </li> </ul>
auth_verbose	Whether or not detailed messages about virtual server authentication should be written to the error log. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
bandwidth_class	The bandwidth management class that this server should use, if any. <ul style="list-style-type: none"> <li>Type: Reference(config-bandwidth)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
completion_rules	Rules that are run at the end of a transaction, in order, comma separated. <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connect_timeout	The time, in seconds, for which an established connection can remain idle waiting for some initial data to be received from the client. The initial data is defined as a complete set of request headers for HTTP, SIP and RTSP services, or the first byte of data for all other services. A value of 0 will disable the timeout. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
connection_errors_error_file	The error message to be sent to the client when the traffic manager detects an internal or backend error for the virtual server. <ul style="list-style-type: none"> <li>Type: Reference(config-extra-file)</li> <li>Required: false</li> <li>Default value: "Default"</li> </ul>

Property	Description
connection_keepalive	<p>Whether or not the virtual server should use keepalive connections with the remote clients.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
connection_keepalive_timeout	<p>The length of time that the virtual server should keep an idle keepalive connection before discarding it. A value of 0 (zero) will mean that the keepalives are never closed by the traffic manager.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
connection_max_client_buffer	<p>The amount of memory, in bytes, that the virtual server should use to store data sent by the client. Larger values will use more memory, but will minimise the number of read() and write() system calls that the traffic manager must perform.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65536"</li> </ul>
connection_max_server_buffer	<p>The amount of memory, in bytes, that the virtual server should use to store data returned by the server. Larger values will use more memory, but will minimise the number of read() and write() system calls that the traffic manager must perform.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65536"</li> </ul>
connection_max_transaction_duration	<p>The total amount of time a transaction can take, counted from the first byte being received until the transaction is complete. For HTTP, this can mean all data has been written in both directions, or the connection has been closed; in most other cases it is the same as the connection being closed. The default value of 0 means there is no maximum duration, i.e., transactions can take arbitrarily long if none of the other timeouts occur.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
connection_server_first_banner	<p>If specified, the traffic manager will use the value as the banner to send for server-first protocols such as FTP, POP, SMTP and IMAP. This allows rules to use the first part of the client data (such as the username) to select a pool. The banner should be in the correct format for the protocol, e.g. for FTP it should start with "220 "</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>



Property	Description
connection_timeout	<p>A connection should be closed if no additional data has been received for this period of time. A value of 0 (zero) will disable this timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "300"</li> </ul>
cookie_domain	<p>The way in which the traffic manager should rewrite the domain portion of any cookies set by a back-end web server.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "no_rewrite"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"no_rewrite": Do not rewrite the domain</li> <li>"set_to_named": Rewrite the domain to the named domain value</li> <li>"set_to_request": Rewrite the domain to the host header of the request</li> </ul> </li> </ul>
cookie_new_domain	<p>The domain to use when rewriting a cookie's domain to a named value.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cookie_path_regex	<p>If you wish to rewrite the path portion of any cookies set by a back-end web server, provide a regular expression to match the path:</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cookie_path_replace	<p>If cookie path regular expression matches, it will be replaced by this substitution. Parameters \$1-\$9 can be used to represent bracketed parts of the regular expression.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
cookie_secure	<p>Whether or not the traffic manager should modify the "secure" tag of any cookies set by a back-end web server.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "no_modify"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"no_modify": Do not modify the 'secure' tag</li> <li>"set_secure": Set the 'secure' tag</li> <li>"unset_secure": Unset the 'secure' tag</li> </ul> </li> </ul>
dns_edns_client_subnet	<p>Enable/Disable use of EDNS client subnet option</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
dns_edns_udpsize	EDNS UDP size advertised in responses. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>
dns_max_udpsize	Maximum UDP answer size. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>
dns_rrset_order	Response record ordering. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "fixed"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"cyclic": Cyclic</li> <li>"fixed": Fixed</li> </ul> </li> </ul>
dns_verbose	Whether or not the DNS Server should emit verbose logging. This is useful for diagnosing problems. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
dns_zones	The DNS zones <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
enabled	Whether the virtual server is enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ftp_data_source_port	The source port to be used for active-mode FTP data connections. If 0, a random high port will be used, otherwise the specified port will be used. If a port below 1024 is required you must first explicitly permit use of low ports with the data_bind_low global setting. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ftp_force_client_secure	Whether or not the virtual server should require that incoming FTP data connections from the client originate from the same IP address as the corresponding client control connection. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>

Property	Description
ftp_force_server_secure	Whether or not the virtual server should require that incoming FTP data connections from the nodes originate from the same IP address as the node. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ftp_port_range_high	If non-zero, then this controls the upper bound of the port range to use for FTP data connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ftp_port_range_low	If non-zero, then this controls the lower bound of the port range to use for FTP data connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ftp_ssl_data	Use SSL on the data connection as well as the control connection (if not enabled it is left to the client and server to negotiate this). <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
glb_services	The associated GLB services for this DNS virtual server. <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
gzip_compress_level	Compression level (1-9, 1=low, 9=high). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>
gzip_enabled	Compress web pages sent back by the server. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
gzip_etag_rewrite	How the ETag header should be manipulated when compressing content. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "wrap"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"delete": Delete the ETag header</li> <li>"ignore": Leave the ETag unchanged</li> <li>"weaken": Change the ETag header to specify a weak match</li> <li>"wrap": Wrap the ETag, and attempt to unwrap safe conditional requests</li> </ul> </li> </ul>

Property	Description
gzip_include_mime	<p>MIME types to compress. Complete MIME types can be used, or a type can end in a '*' to match multiple types.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: "text/html text/plain"</li> </ul>
gzip_max_size	<p>Maximum document size to compress (0 means unlimited).</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10000000"</li> </ul>
gzip_min_size	<p>Minimum document size to compress.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1000"</li> </ul>
gzip_no_size	<p>Compress documents with no given size.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http2_connect_timeout	<p>The time, in seconds, to wait for a request on a new HTTP/2 connection. If no request is received within this time, the connection will be closed. This setting overrides the connect_timeout setting. If set to 0 (zero), the value of connect_timeout will be used instead.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http2_data_frame_size	<p>This setting controls the preferred frame size used when sending body data to the client. If the client specifies a smaller maximum size than this setting, the client's maximum size will be used. Every data frame sent has at least a 9-byte header, in addition to this frame size, prepended to it.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>
http2_enabled	<p>This setting allows the HTTP/2 protocol to be used by a HTTP virtual server. Unless use of HTTP/2 is negotiated by the client, the virtual server will fall back to HTTP 1.x automatically.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http2_header_table_size	<p>This setting controls the amount of memory allowed for header compression on each HTTP/2 connection.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "4096"</li> </ul>

Property	Description
http2_headers_index_blacklist	<p>A list of header names that should never be compressed using indexing.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http2_headers_index_default	<p>The HTTP/2 HPACK compression scheme allows for HTTP headers to be compressed using indexing. Sensitive headers can be marked as "never index", which prevents them from being compressed using indexing. When this setting is Yes, only headers included in http2!headers_index_blacklist are marked as "never index". When this setting is No, all headers will be marked as "never index" unless they are included in http2!headers_index_whitelist.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http2_headers_index_whitelist	<p>A list of header names that can be compressed using indexing when the value of http2!headers_index_default is set to No.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http2_headers_size_limit	<p>The maximum size, in bytes, of decompressed headers for an HTTP/2 request. If the limit is exceeded, the connection on which the request was sent will be dropped. A value of 0 disables the limit check. If a service protection class with http!max_header_length configured is associated with this service then that setting will take precedence.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "262144"</li> </ul>
http2_idle_timeout_no_streams	<p>The time, in seconds, to wait for a new HTTP/2 request on a previously used HTTP/2 connection that has no open HTTP/2 streams. If an HTTP/2 request is not received within this time, the connection will be closed. A value of 0 (zero) will disable the timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "120"</li> </ul>
http2_idle_timeout_open_streams	<p>The time, in seconds, to wait for data on an idle HTTP/2 connection, which has open streams, when no data has been sent recently (e.g. for long-poll requests). If data is not sent within this time, all open streams and the HTTP/2 connection will be closed. A value of 0 (zero) will disable the timeout.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "600"</li> </ul>

Property	Description
http2_max_concurrent_streams	<p>This setting controls the number of streams a client is permitted to open concurrently on a single connection.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "200"</li> </ul>
http2_max_frame_size	<p>This setting controls the maximum HTTP/2 frame size clients are permitted to send to the traffic manager.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "16384"</li> </ul>
http2_max_header_padding	<p>The maximum size, in bytes, of the random-length padding to add to HTTP/2 header frames. The padding, a random number of zero bytes up to the maximum specified.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http2_merge_cookie_headers	<p>Whether Cookie headers received from an HTTP/2 client should be merged into a single Cookie header using RFC6265 rules before forwarding to an HTTP/1.1 server. Some web applications do not handle multiple Cookie headers correctly.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http2_stream_window_size	<p>This setting controls the flow control window for each HTTP/2 stream. This will limit the memory used for buffering when the client is sending body data faster than the pool node is reading it.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65535"</li> </ul>
http_add_cluster_ip	<p>Whether or not the virtual server should add an "X-Cluster-Client-Ip" header to the request that contains the remote client's IP address.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http_add_x_forwarded_for	<p>Whether or not the virtual server should append the remote client's IP address to the X-Forwarded-For header. If the header does not exist, it will be added.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
http_add_x_forwarded_proto	<p>Whether or not the virtual server should add an "X-Forwarded-Proto" header to the request that contains the original protocol used by the client to connect to the traffic manager.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
http_autodetect_upgrade_headers	<p>Whether the traffic manager should check for HTTP responses that confirm an HTTP connection is transitioning to the WebSockets protocol. If that such a response is detected, the traffic manager will cease any protocol-specific processing on the connection and just pass incoming data to the client/server as appropriate.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
http_chunk_overhead_forwarding	<p>Handling of HTTP chunk overhead. When vTM receives data from a server or client that consists purely of protocol overhead (contains no payload), forwarding of such segments is delayed until useful payload data arrives (setting "lazy"). Changing this key to "eager" will make vTM incur the overhead of immediately passing such data on; it should only be used with HTTP peers whose chunk handling requires it.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "lazy"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"eager": Forward all data, even when no new payload information is available.</li> <li>"lazy": Only forward segments when useful payload data is available.</li> </ul> </li> </ul>
http_location_regex	<p>If the 'Location' header matches this regular expression, rewrite the header using the 'location_replace' pattern.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
http_location_replace	<p>If the 'Location' header matches the 'location_regex' regular expression, rewrite the header with this pattern (parameters such as \$1-\$9 can be used to match parts of the regular expression):</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
http_location_rewrite	<p>The action the virtual server should take if the "Location" header does not match the location_regex regular expression.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "if_host_matches"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"always": Rewrite the hostname to the request's "Host" header, and rewrite the protocol and port if necessary;</li> <li>"if_host_matches": Do not rewrite the hostname. Rewrite the protocol and port if the hostname matches the request's "Host" header.</li> <li>"never": Nothing;</li> </ul> </li> </ul>
http_mime_default	<p>Auto-correct MIME types if the server sends the "default" MIME type for files.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "text/plain"</li> </ul>
http_mime_detect	<p>Auto-detect MIME types if the server does not provide them.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
http_strip_x_forwarded_proto	<p>Whether or not the virtual server should strip the 'X-Forwarded-Proto' header from incoming requests.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
kerberos_protocol_transition_enabled	<p>Whether or not the virtual server should use Kerberos Protocol Transition.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
kerberos_protocol_transition_principal	<p>The Kerberos principal this virtual server should use to perform Kerberos Protocol Transition.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
kerberos_protocol_transition_target	<p>The Kerberos principal name of the service this virtual server targets.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
listen_on_any	<p>Whether to listen on all IP addresses</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>



Property	Description
listen_on_hosts	Hostnames and IP addresses to listen on <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
listen_on_traffic_ips	Traffic IP Groups to listen on <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
log_client_connection_failures	Should the virtual server log failures occurring on connections to clients. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
log_enabled	Whether or not to log connections to the virtual server to a disk on the file system. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
log_filename	The name of the file in which to store the request logs. The filename can contain macros which will be expanded by the traffic manager to generate the full filename. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "%zeushome%/zxtm/log/%v.log"</li> </ul>
log_format	The log file format. This specifies the line of text that will be written to the log file when a connection to the traffic manager is completed. Many parameters from the connection can be recorded using macros. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "%h %l %u %t "%r" %s %b "%{Referer}i" "%{User-agent}i""</li> </ul>
log_save_all	Whether to log all connections by default, or log no connections by default. Specific connections can be selected for addition to or exclusion from the log using the TrafficScript function requestlog.include(). <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
log_server_connection_failures	Should the virtual server log failures occurring on connections to nodes. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
log_session_persistence_verbose	Should the virtual server log session persistence events. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
log_ssl_failures	Should the virtual server log failures occurring on SSL secure negotiation. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
log_ssl_resumption_failures	Should the virtual server log messages when attempts to resume SSL sessions (either from the session cache or a session ticket) fail. Note that failure to resume an SSL session does not result in the SSL connection being closed, but it does cause a full SSL handshake to take place. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
max_concurrent_connections	The maximum number of concurrent TCP connections that will be handled by this virtual server. If set to a non-zero value, the traffic manager will limit the number of concurrent TCP connections that this virtual server will accept to the value specified. When the limit is reached, new connections to this virtual server will not be accepted. If set to 0 the number of concurrent TCP connections will not be limited. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
note	A description for the virtual server. <ul style="list-style-type: none"> <li>Type: FreeformString</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
pool	The default pool to use for traffic. <ul style="list-style-type: none"> <li>Type: Reference(config-pool)</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
port	The port on which to listen for incoming connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: true</li> <li>Default value: &lt;none&gt;</li> </ul>
protection_class	The service protection class that should be used to protect this server, if any. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
protocol	<p>The protocol that the virtual server is using.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "http"</li> <li>• Permitted values: <ul style="list-style-type: none"> <li>"client_first": Generic client first</li> <li>"dns": DNS (UDP)</li> <li>"dns_tcp": DNS (TCP)</li> <li>"ftp": FTP</li> <li>"http": HTTP</li> <li>"https": SSL (HTTPS)</li> <li>"imaps": SSL (IMAPS)</li> <li>"imapv2": IMAPv2</li> <li>"imapv3": IMAPv3</li> <li>"imapv4": IMAPv4</li> <li>"ldap": LDAP</li> <li>"ldaps": SSL (LDAPS)</li> <li>"pop3": POP3</li> <li>"pop3s": SSL (POP3S)</li> <li>"rtsp": RTSP</li> <li>"server_first": Generic server first</li> <li>"siptcp": SIP (TCP)</li> <li>"sipudp": SIP (UDP)</li> <li>"smtp": SMTP</li> <li>"ssl": SSL</li> <li>"stream": Generic streaming</li> <li>"telnet": Telnet</li> <li>"udp": UDP</li> <li>"udpstreaming": UDP - Streaming</li> </ul> </li> </ul>
proxy_protocol	<p>Expect connections to the traffic manager to be prefixed with a PROXY protocol header. If enabled, the information contained in the PROXY header will be available in TrafficScript. Connections that are not prefixed with a valid PROXY protocol header will be discarded.</p> <ul style="list-style-type: none"> <li>• Type: Boolean</li> <li>• Required: false</li> <li>• Default value: false</li> </ul>

Property	Description
recent_connections_enabled	Whether or not connections handled by this virtual server should be shown on the Activity > Connections page. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
recent_connections_save_all	Whether or not all connections handled by this virtual server should be shown on the Connections page. Individual connections can be selectively shown on the Connections page using the recentconns.include() TrafficScript function. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
request_rules	Rules to be applied to incoming requests, in order, comma separated. <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
request_tracing_enabled	Record a trace of major connection processing events for each request and response. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
request_tracing_trace_io	Include details of individual I/O events in request and response traces. Requires request tracing to be enabled. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
response_rules	Rules to be applied to responses, in order, comma separated. <ul style="list-style-type: none"> <li>Type: List(Reference(config-trafficscript))</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rtsp_streaming_port_range_high	If non-zero this controls the upper bound of the port range to use for streaming data connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
rtsp_streaming_port_range_low	If non-zero this controls the lower bound of the port range to use for streaming data connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
rtsp_streaming_timeout	<p>If non-zero data-streams associated with RTSP connections will timeout if no data is transmitted for this many seconds.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
sip_dangerous_requests	<p>The action to take when a SIP request with body data arrives that should be routed to an external IP.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "node"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"forbid": Send a 403 Forbidden response to the client</li> <li>"forward": Forward the request to its target URI (dangerous)</li> <li>"node": Send the request to a back-end node</li> </ul> </li> </ul>
sip_follow_route	<p>Should the virtual server follow routing information contained in SIP requests. If set to No requests will be routed to the chosen back-end node regardless of their URI or Route header.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
sip_max_connection_mem	<p>SIP clients can have several pending requests at one time. To protect the traffic manager against DoS attacks, this setting limits the amount of memory each client can use. When the limit is reached new requests will be sent a 413 response. If the value is set to 0 (zero) the memory limit is disabled.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "65536"</li> </ul>
sip_mode	<p>The mode that this SIP virtual server should operate in.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "sip_gateway"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"full_gateway": All SIP requests and responses and all session data will pass through vTM. A port range to use for the session data and a timeout value for inactive data connections can be specified in the additional settings that are displayed when the Full Gateway mode is selected.</li> <li>"route": The first SIP request in a session will pass through vTM, along with its responses, but all future requests that are part of the same session will go directly to the back-end node that was chosen by the traffic manager.</li> <li>"sip_gateway": All SIP requests and responses will pass through the traffic manager.</li> </ul> </li> </ul>

Property	Description
sip_rewrite_uri	<p>Replace the Request-URI of SIP requests with the address of the selected back-end node.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
sip_streaming_port_range_high	<p>If non-zero this controls the upper bound of the port range to use for streaming data connections.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sip_streaming_port_range_low	<p>If non-zero, then this controls the lower bound of the port range to use for streaming data connections.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
sip_streaming_timeout	<p>If non-zero a UDP stream will timeout when no data has been seen within this time.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "60"</li> </ul>
sip_timeout_messages	<p>When timing out a SIP transaction, send a 'timed out' response to the client and, in the case of an INVITE transaction, a CANCEL request to the server.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
sip_transaction_timeout	<p>The virtual server should discard a SIP transaction when no further messages have been seen within this time.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
sip_udp_associate_by_source	<p>Require that SIP datagrams which are part of the same transaction are received from the same address and port.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
slm_class	<p>The service level monitoring class that this server should use, if any.</p> <ul style="list-style-type: none"> <li>Type: Reference(config-slm)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
smtp_expect_starttls	<p>Whether or not the traffic manager should expect the connection to start off in plain text and then upgrade to SSL using STARTTLS when handling SMTP traffic.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_add_http_headers	<p>Whether or not the virtual server should add HTTP headers to each request to show the SSL connection parameters.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_cipher_suites	<p>The SSL/TLS cipher suites to allow for connections to this virtual server. Leaving this empty will make the virtual server use the globally configured cipher suites, see configuration key <code>ssl!cipher_suites</code> in the Global Settings section of the System tab. See there for how to specify SSL/TLS cipher suites.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_client_cert_cas	<p>The certificate authorities that this virtual server should trust to validate client certificates. If no certificate authorities are selected, and client certificates are requested, then all client certificates will be accepted.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_client_cert_headers	<p>What HTTP headers the virtual server should add to each request to show the data in the client certificate.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "none"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"all": Certificate fields and certificate text</li> <li>"none": No data</li> <li>"simple": Certificate fields</li> </ul> </li> </ul>
ssl_decrypt	<p>Whether or not the virtual server should decrypt incoming SSL traffic.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
ssl_elliptic_curves	<p>The SSL elliptic curve preference list for SSL connections to this virtual server using TLS version 1.0 or higher. Leaving this empty will make the virtual server use the globally configured curve preference list. The named curves P256, P384 and P521 may be configured.</p> <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_honor_fallback_scsv	<p>Whether or not the Fallback SCSV sent by TLS clients is honored by this virtual server. Choosing the global setting means the value of configuration key <code>ssl!honor_fallback_scsv</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable Fallback SCSV</li> <li>"enabled": Enable Fallback SCSV</li> <li>"use_default": Use the global setting for Fallback SCSV</li> </ul> </li> </ul>
ssl_issued_certs_never_expire	<p>When the virtual server verifies certificates signed by these certificate authorities, it doesn't check the 'not after' date, i.e., they are considered valid even after their expiration date has passed (but not if they have been revoked).</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_issued_certs_never_expire_depth	<p>This setting gives the number of certificates in a certificate chain beyond those listed as <code>issued_certs_never_expire</code> whose certificate expiry will not be checked. For example "0" will result in the expiry checks being made for certificates issued by <code>issued_certs_never_expire</code> certificates, "1" will result in no expiry checks being performed for the certificates directly issued by <code>issued_certs_never_expire</code> certificates, "2" will avoid checking expiry for certificates issued by certificates issued by the <code>issued_certs_never_expire</code> certificates as well, and so on.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>
ssl_ocsp_enable	<p>Whether or not the traffic manager should use OCSP to check the revocation status of client certificates.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>



Property	Description
ssl_ocsp_issuers	<p>A table of certificate issuer specific OCSP settings.</p> <ul style="list-style-type: none"> <li>• Type: Table</li> <li>• Required: false</li> <li>• Primary key: <ul style="list-style-type: none"> <li>- issuer (String): The name of an issuer (or DEFAULT for default OCSP settings). (Required)</li> </ul> </li> <li>• Sub keys: <ul style="list-style-type: none"> <li>- aia (Boolean): Whether the traffic manager should use AIA information contained in a client certificate to determine which OCSP responder to contact.</li> <li>- nonce (Enum(String)): How to use the OCSP nonce extension, which protects against OCSP replay attacks. Some OCSP servers do not support nonces.</li> </ul> </li> </ul> <p>Permitted values:</p> <p>"off": No nonce check</p> <p>"on": Use nonce, server does not have to reply with nonce</p> <p>"strict": Use nonce, server must reply with nonce</p> <ul style="list-style-type: none"> <li>- required (Enum(String)): Whether we should do an OCSP check for this issuer, and whether it is required or optional.</li> </ul> <p>Permitted values:</p> <p>"none": None</p> <p>"optional": OCSP check optional</p> <p>"strict": OCSP check required</p> <ul style="list-style-type: none"> <li>- responder_cert (String): The expected responder certificate.</li> <li>- signer (String): The certificate with which to sign the request, if any.</li> <li>- url (String): Which OCSP responders this virtual server should use to verify client certificates.</li> </ul>
ssl_ocsp_issuers_json	<p>A Traffic Manager REST-compatible JSON representation of the "ssl_ocsp_issuers" table property. Use this field with the "vtm_virtual_server_ocsp_issuers_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ssl_ocsp_max_response_age	<p>The number of seconds for which an OCSP response is considered valid if it has not yet exceeded the time specified in the 'nextUpdate' field. If set to 0 (zero) then OCSP responses are considered valid until the time specified in their 'nextUpdate' field.</p> <ul style="list-style-type: none"> <li>• Type: UInt</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>

Property	Description
ssl_ocsp_stapling	<p>If OCSP URIs are present in certificates used by this virtual server, then enabling this option will allow the traffic manager to provide OCSP responses for these certificates as part of the handshake, if the client sends a TLS status_request extension in the ClientHello.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
ssl_ocsp_time_tolerance	<p>The number of seconds outside the permitted range for which the 'thisUpdate' and 'nextUpdate' fields of an OCSP response are still considered valid.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>
ssl_ocsp_timeout	<p>The number of seconds after which OCSP requests will be timed out.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "10"</li> </ul>
ssl_request_client_cert	<p>Whether or not the virtual server should request an identifying SSL certificate from each client.</p> <ul style="list-style-type: none"> <li>Type: Enum(UInt)</li> <li>Required: false</li> <li>Default value: "dont_request"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"dont_request": Do not request a client certificate</li> <li>"request": Request, but do not require a client certificate</li> <li>"require": Require a client certificate</li> </ul> </li> </ul>
ssl_send_close_alerts	<p>Whether or not to send an SSL/TLS "close alert" when the traffic manager is initiating an SSL socket disconnection.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
ssl_server_cert_alt_certificates	<p>The SSL certificates and corresponding private keys.</p> <ul style="list-style-type: none"> <li>Type: List(String)</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_server_cert_default	<p>The default SSL certificate to use for this virtual server.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
ssl_server_cert_host_mapping	<p>Host specific SSL server certificate mappings.</p> <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>host (String): Host which this entry refers to. (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>certificate (String): The SSL server certificate for a particular destination site IP. (Required)</li> <li>alt_certificates (List(String)): The SSL server certificates for a particular destination site IP.</li> </ul> </li> </ul>
ssl_server_cert_host_mapping_json	<p>A Traffic Manager REST-compatible JSON representation of the "ssl_server_cert_host_mapping" table property. Use this field with the "vtm_virtual_server_server_cert_host_mapping_table" data source for dynamic table generation from input variables.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
ssl_session_cache_enabled	<p>Whether or not use of the session cache is enabled for this virtual server. Choosing the global setting means the value of configuration key ssl!session_cache_enabled from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable use of the session cache</li> <li>"enabled": Enable use of the session cache</li> <li>"use_default": Use the global setting for use of the session cache</li> </ul> </li> </ul>
ssl_session_tickets_enabled	<p>Whether or not use of session tickets is enabled for this virtual server. Choosing the global setting means the value of configuration key ssl!tickets!enabled from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable use of the session tickets</li> <li>"enabled": Enable use of the session tickets</li> <li>"use_default": Use the global setting for use of session tickets</li> </ul> </li> </ul>

Property	Description
ssl_signature_algorithms	<p>The SSL signature algorithms preference list for SSL connections to this virtual server using TLS version 1.2 or higher. Leaving this empty will make the virtual server use the globally configured preference list, signature_algorithms in the ssl section of the global_settings resource. See there and in the online help for how to specify SSL signature algorithms.</p> <ul style="list-style-type: none"> <li>• Type: String</li> <li>• Required: false</li> <li>• Default value: &lt;none&gt;</li> </ul>
ssl_support_ssl3	<p>Whether or not SSLv3 is enabled for this virtual server. Choosing the global setting means the value of configuration key sslsupport_ssl3 from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "use_default"</li> <li>• Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable SSLv3</li> <li>"enabled": Enable SSLv3</li> <li>"use_default": Use the global setting for SSLv3</li> </ul> </li> </ul>
ssl_support_tls1	<p>Whether or not TLSv1.0 is enabled for this virtual server. Choosing the global setting means the value of configuration key sslsupport_tls1 from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "use_default"</li> <li>• Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.0</li> <li>"enabled": Enable TLSv1.0</li> <li>"use_default": Use the global setting for TLSv1.0</li> </ul> </li> </ul>
ssl_support_tls1_1	<p>Whether or not TLSv1.1 is enabled for this virtual server. Choosing the global setting means the value of configuration key sslsupport_tls1_1 from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>• Type: Enum(String)</li> <li>• Required: false</li> <li>• Default value: "use_default"</li> <li>• Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.1</li> <li>"enabled": Enable TLSv1.1</li> <li>"use_default": Use the global setting for TLSv1.1</li> </ul> </li> </ul>

Property	Description
ssl_support_tls1_2	<p>Whether or not TLSv1.2 is enabled for this virtual server. Choosing the global setting means the value of configuration key <code>ssl!support_tls1_2</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.2</li> <li>"enabled": Enable TLSv1.2</li> <li>"use_default": Use the global setting for TLSv1.2</li> </ul> </li> </ul>
ssl_support_tls1_3	<p>Whether or not TLSv1.3 is enabled for this virtual server. Choosing the global setting means the value of configuration key <code>ssl!support_tls1_3</code> from the Global Settings section of the System tab will be enforced.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "use_default"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"disabled": Disable TLSv1.3</li> <li>"enabled": Enable TLSv1.3</li> <li>"use_default": Use the global setting for TLSv1.3</li> </ul> </li> </ul>
ssl_trust_magic	<p>If the traffic manager is receiving traffic sent from another traffic manager, then enabling this option will allow it to decode extra information on the true origin of the SSL connection. This information is supplied by the first traffic manager.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
syslog_enabled	<p>Whether or not to log connections to the virtual server to a remote syslog host.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
syslog_format	<p>The log format for the remote syslog. This specifies the line of text that will be sent to the remote syslog when a connection to the traffic manager is completed. Many parameters from the connection can be recorded using macros.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: "%h %l %u %t \"%r\" %s %b \"%{Referer}i\" \"%{User-agent}i\""</li> </ul>
syslog_ip_end_point	<p>The remote host and port (default is 514) to send request log lines to.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

Property	Description
syslog_msg_len_limit	<p>Maximum length in bytes of a message sent to the remote syslog. Messages longer than this will be truncated before they are sent.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2048"</li> </ul>
tcp_close_with_rst	<p>Whether or not connections from clients should be closed with a RST packet, rather than a FIN packet. This avoids the TIME_WAIT state, which on rare occasions allows wandering duplicate packets to be safely ignored.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
tcp_nagle	<p>Whether or not Nagle's algorithm should be used for TCP connections.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
tcp_proxy_close	<p>If set to Yes the traffic manager will send the client FIN to the back-end server and wait for a server response instead of closing the connection immediately. This is only necessary for protocols that require half-close support to function correctly, such as "rsh". If the traffic manager is responding to the request itself, setting this key to Yes will cause the traffic manager to continue writing the response even after it has received a FIN from the client.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
transaction_export_brief	<p>Whether to export a restricted set of metadata about transactions processed by this virtual server. If enabled, more verbose information such as client and server headers and request tracing events will be omitted from the exported data.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
transaction_export_enabled	<p>Export metadata about transactions handled by this service to the globally configured endpoint. Data will be exported only if the global transaction_export!enabled setting is enabled.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
transaction_export_hi_res	<p>Whether the transaction processing timeline included in the metadata export is recorded with a high, microsecond, resolution. If set to No, timestamps will be recorded with a resolution of milliseconds.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
transaction_export_http_header_blacklist	<p>The set of HTTP header names for which corresponding values should be redacted from the metadata exported by this virtual server.</p> <ul style="list-style-type: none"> <li>Type: Set(String)</li> <li>Required: false</li> <li>Default value: "Authorization"</li> </ul>
transparent	<p>Whether or not bound sockets should be configured for transparent proxying.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
udp_end_point_persistence	<p>Whether UDP datagrams received from the same IP address and port are sent to the same pool node if they match an existing UDP session. Sessions are defined by the protocol being handled, for example SIP datagrams are grouped based on the value of the Call-ID header.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: true</li> </ul>
udp_port_smp	<p>Whether or not UDP datagrams should be distributed across all traffic manager processes, if this behaviour is not normally selected automatically due to other settings.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
udp_rbuff_size	<p>If this setting is non-zero, the virtual server will set the socket receive buffer size to this number of bytes. If set, this will override the so_rbuff_size setting. An OS-specified limit on socket buffer sizes such as given by sysctl net.core.rmem_max can be exceeded using this setting.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
udp_response_datagrams_expired	<p>The virtual server should discard any UDP connection and reclaim resources when the node has responded with this number of datagrams. For simple request/response protocols this can be often set to 1. If set to -1, the connection will not be discarded until the timeout is reached.</p> <ul style="list-style-type: none"> <li>Type: Int</li> <li>Required: false</li> <li>Default value: "1"</li> </ul>

Property	Description
udp_smp_mode	<p>Whether the traffic manager should try to use SO_REUSEPORT for distributing incoming UDP datagrams across multiple processes (if kernel support is detected) or whether the legacy (pre-20.2) multi-processing mode should be used.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "auto"</li> <li>Permitted values: "auto": auto "legacy": legacy</li> </ul>
udp_timeout	<p>The virtual server should discard any UDP connection and reclaim resources when no further UDP traffic has been seen within this time.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "7"</li> </ul>
udp_wbuff_size	<p>If this setting is non-zero, the virtual server will set the socket send buffer size to this number of bytes. If set, this will override the so_wbuff_size setting. An OS-specified limit on socket buffer sizes such as given by sysctl net.core.wmem_max can be exceeded using this setting.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
web_cache_control_out	<p>The "Cache-Control" header to add to every cached HTTP response, no-cache or max-age=600 for example.</p> <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
web_cache_enabled	<p>If set to Yes the traffic manager will attempt to cache web server responses.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>
web_cache_error_page_time	<p>Time period to cache error pages for.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "30"</li> </ul>



Property	Description
web_cache_max_time	Maximum time period to cache web pages for. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "600"</li> </ul>
web_cache_refresh_time	If a cached page is about to expire within this time, the traffic manager will start to forward some new requests on to the web servers. A maximum of one request per second will be forwarded; the remainder will continue to be served from the cache. This prevents "bursts" of traffic to your web servers when an item expires from the cache. Setting this value to 0 will stop the traffic manager updating the cache before it expires. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: "2"</li> </ul>

### vtm\_optimizer\_profile

URI Endpoint: /api/tm/8.1/config/active/optimizer/profiles

A Web Accelerator profile can be applied to an HTTP virtual server to enable automatic web content optimization.

```
# Example usage of vtm_optimizer_profile resource...
resource "vtm_optimizer_profile" "my_optimizer_profile" {
  name = "MyOptimizerProfile"
}
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
background_after	If Web Accelerator can finish optimizing the resource within this time limit then serve the optimized content to the client, otherwise complete the optimization in the background and return the original content to the client. If set to 0, Web Accelerator will always wait for the optimization to complete before sending a response to the client. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
background_on_additional_resources	If a web page contains resources that have not yet been optimized, fetch and optimize those resources in the background and send a partially optimized web page to clients until all resources on that page are ready. <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

Property	Description
mode	<p>Set the Web Accelerator mode to turn acceleration on or off.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Required: false</li> <li>Default value: "active"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"active": On - Web Accelerator acceleration is enabled</li> <li>"idle": Off - Acceleration is disabled, but requests for Web Accelerator resources are served</li> <li>"stealth": Stealth - Acceleration is controlled by a cookie</li> </ul> </li> </ul>
show_info_bar	<p>Show the Web Accelerator information bar on optimized web pages. This requires HTML optimization to be enabled in the acceleration settings.</p> <ul style="list-style-type: none"> <li>Type: Boolean</li> <li>Required: false</li> <li>Default value: false</li> </ul>

## Data Sources

Data sources are used to obtain information about an object or environment.

## Resources

Each configuration resource is also available as an identical Terraform "data source" object. All properties in a data source remain the same as the configuration resource on which it is based, but read-only. Configuration data sources allow a Terraform template to query the contents of a Traffic Manager configuration object that was created externally from the template. For example:

```
# Create a data source object to represent an existing virtual server
data "vtm_virtual_server" "my_vs_data" {
  name = "MyVS"
}

# Set up a new virtual server referencing the data source
resource "vtm_virtual_server" "my_new_vs" {
  name = "MyNewVS"
  pool = "discard"
  port = 1234
  ssl_decrypt = true
  ssl_server_cert_default =
"${data.vtm_virtual_server.my_vs_data.ssl_server_cert_default}"
}
```

For details of the available properties for a particular resource, refer to the corresponding configuration resource reference earlier in this chapter.

## Resource Lists

Resource list data sources represent a list of stored configuration objects of the named type. For example, to obtain the current list of virtual servers, use the "vtm\_virtual\_server\_list" data source. By default, a resource list data source returns all objects of the given type. To filter the results, use the "starts\_with", "ends\_with", "contains", or "regex\_match" filters in the data source definition. To obtain the resulting filtered object list, use the "object\_list" property.

```
# Create a list of all pools
data "vtm_pool_list" "my_pool_list" {}

# Create a list of virtual servers with names starting "Exchange_"
data "vtm_virtual_server_list" "my_vs_list" {
  starts_with = "Exchange_"
}

# Create a TrafficScript rule list based on a regex match
data "vtm_rule_list" "my_rule_list" {
  regex_match = "^Rule-.*?-[0-9]$"
}

# Use the rule list in a virtual server definition
resource "vtm_virtual_server" "my_vs" {
  ...
  request_rules = ["${data.vtm_rule_list.my_rule_list.object_list}"]
  ...
}
```

The available resource list data sources are:

- vtm\_action\_list
- vtm\_action\_program\_list
- vtm\_optimizer\_profile\_list
- vtm\_optimizer\_scope\_list
- vtm\_bandwidth\_list
- vtm\_bgpneighbor\_list
- vtm\_cloud\_api\_credential\_list
- vtm\_custom\_list
- vtm\_dns\_server\_zone\_file\_list
- vtm\_dns\_server\_zone\_list
- vtm\_event\_type\_list
- vtm\_extra\_file\_list
- vtm\_glb\_service\_list
- vtm\_kerberos\_keytab\_list
- vtm\_kerberos\_krb5conf\_list

- vtm\_kerberos\_principal\_list
- vtm\_license\_key\_list
- vtm\_location\_list
- vtm\_log\_export\_list
- vtm\_monitor\_list
- vtm\_monitor\_script\_list
- vtm\_persistence\_list
- vtm\_pool\_list
- vtm\_protection\_list
- vtm\_rate\_list
- vtm\_rule\_authenticator\_list
- vtm\_rule\_list
- vtm\_saml\_trustedidp\_list
- vtm\_service\_level\_monitor\_list
- vtm\_servicediscovery\_list
- vtm\_ssl\_admin\_ca\_list
- vtm\_ssl\_ca\_list
- vtm\_ssl\_client\_key\_list
- vtm\_ssl\_server\_key\_list
- vtm\_ssl\_ticket\_key\_list
- vtm\_traffic\_ip\_group\_list
- vtm\_traffic\_manager\_list
- vtm\_user\_authenticator\_list
- vtm\_user\_group\_list
- vtm\_virtual\_server\_list

## Statistics

Statistics data sources allow Terraform to access a Traffic Manager's SNMP counter values. Some statistics data sources access single objects and take no parameters (for example, vtm\_globals\_stats), while others represent a collection of objects and require a "name" parameter (for example, vtm\_virtual\_server\_stats):

```
# Single object
data "vtm_globals_stats" "global_stats" {}

# Collection of objects
data "vtm_virtual_server_stats" "vs_stats" {
```

```

    name = "MyVirtualServer"
  }

```

### vtm\_action\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/actions/\*

Actions statistics values.

```

# Example usage of vtm_action_stats data source...
data "vtm_action_stats" "my_action_stats" {
  name = "MyAction"
}

# Access with...
"${data.vtm_action_stats.my_action_stats.<FIELD_NAME>}"

```

Field Name	Description
processed	Number of times this action has been processed. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "actionsProcessed"</li> </ul>

### vtm\_cache\_asp\_session\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/asp\_session\_cache

Asp session cache statistics values.

```

# Example usage of vtm_cache_asp_session_cache_stats data source...
data "vtm_cache_asp_session_cache_stats" "my_cache_asp_session_cache_stats" {}

# Access with...
"${data.vtm_cache_asp_session_cache_stats.my_cache_asp_session_cache_stats.<FIELD_NAME>}"

```

Field Name	Description
entries	The total number of ASP sessions stored in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheEntries"</li> </ul>
entries_max	The maximum number of ASP sessions in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheEntriesMax"</li> </ul>
hit_rate	The percentage of ASP session lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheHitRate"</li> </ul>

Field Name	Description
hits	Number of times a ASP session entry has been successfully found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheHits"</li> </ul>
lookups	Number of times a ASP session entry has been looked up in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheLookups"</li> </ul>
misses	Number of times a ASP session entry has not been available in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheMisses"</li> </ul>
oldest	The age of the oldest ASP session in the cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "aspSessionCacheOldest"</li> </ul>

### vtm\_bandwidth\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/bandwidth/\*

Bandwidth statistics values.

```
# Example usage of vtm_bandwidth_stats data source...
data "vtm_bandwidth_stats" "my_bandwidth_stats" {
  name = "MyBandwidth"
}

# Access with...
"${data.vtm_bandwidth_stats.my_bandwidth_stats.<FIELD_NAME>}"
```

Field Name	Description
bytes_out	Bytes output by connections assigned to this bandwidth class. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "bandwidthClassBytesOut"</li> </ul>
guarantee	Guaranteed bandwidth class limit (kbits/s). Currently unused. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "bandwidthClassGuarantee"</li> </ul>
maximum	Maximum bandwidth class limit (kbits/s). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "bandwidthClassMaximum"</li> </ul>

### vtm\_cloud\_api\_credential\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cloud\_api\_credentials/\*

Cloud api credentials statistics values.

```
# Example usage of vtm_cloud_api_credential_stats data source...
data "vtm_cloud_api_credential_stats" "my_cloud_api_credential_stats" {
  name = "MyCloudApiCredential"
}

# Access with...
"${data.vtm_cloud_api_credential_stats.my_cloud_api_credential_stats.<FIELD_NAME>}"
```

Field Name	Description
node_creations	The number of instance creation API requests made with this set of cloud credentials. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "cloudcredentialsNodeCreations"</li> </ul>
node_deletions	The number of instance destruction API requests made with this set of cloud credentials. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "cloudcredentialsNodeDeletions"</li> </ul>
status_requests	The number of status API requests made with this set of cloud credentials. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "cloudcredentialsStatusRequests"</li> </ul>

### vtm\_connection\_rate\_limit\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/connection\_rate\_limit/\*

Connection rate limit statistics values.

```
# Example usage of vtm_connection_rate_limit_stats data source...
data "vtm_connection_rate_limit_stats" "my_connection_rate_limit_stats" {
  name = "MyConnectionRateLimit"
}

# Access with...
"${data.vtm_connection_rate_limit_stats.my_connection_rate_limit_stats.<FIELD_NAME>}"
```

Field Name	Description
conns_entered	Connections that have entered the rate class and have been queued. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassConnsEntered"</li> </ul>
conns_left	Connections that have left the rate class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassConnsLeft"</li> </ul>

Field Name	Description
current_rate	The average rate that requests are passing through this rate class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassCurrentRate"</li> </ul>
dropped	Requests dropped from this rate class without being processed (e.g. timeouts). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassDropped"</li> </ul>
max_rate_per_min	The maximum rate that requests may pass through this rate class (requests/min). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassMaxRatePerMin"</li> </ul>
max_rate_per_sec	The maximum rate that requests may pass through this rate class (requests/sec). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassMaxRatePerSec"</li> </ul>
queue_length	The current number of requests queued by this rate class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rateClassQueueLength"</li> </ul>

### vtm\_event\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/events/\*

Events statistics values.

```
# Example usage of vtm_event_stats data source...
data "vtm_event_stats" "my_event_stats" {
  name = "MyEvent"
}

# Access with...
"${data.vtm_event_stats.my_event_stats.<FIELD_NAME>}"
```

Field Name	Description
matched	Number of times this event configuration has matched. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "eventsMatched"</li> </ul>

### vtm\_glb\_service\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/glb\_services/\*

Glb services statistics values.



```
# Example usage of vtm_glb_service_stats data source...
data "vtm_glb_service_stats" "my_glb_service_stats" {
  name = "MyGlbService"
}

# Access with...
"${data.vtm_glb_service_stats.my_glb_service_stats.<FIELD_NAME>}"
```

Field Name	Description
discarded	Number of A records this GLB Service has discarded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "glbServiceDiscarded"</li> </ul>
responses	Number of A records this GLB Service has altered. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "glbServiceResponses"</li> </ul>
unmodified	Number of A records this GLB Service has passed through unmodified. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "glbServiceUnmodified"</li> </ul>

### vtm\_globals\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/globals

Globals statistics values.

```
# Example usage of vtm_globals_stats data source...
data "vtm_globals_stats" "my_globals_stats" {}

# Access with...
"${data.vtm_globals_stats.my_globals_stats.<FIELD_NAME>}"
```

Field Name	Description
analytics_transactions_dropped	Count of transaction metadata records that have been dropped <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "analyticsTransactionsDropped"</li> </ul>
analytics_transactions_exported	Count of transaction metadata records that have been exported <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "analyticsTransactionsExported"</li> </ul>
analytics_transactions_memory_usage	Number of bytes queued in the transaction export transmit buffers. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "analyticsTransactionsMemoryUsage"</li> </ul>
data_entries	Number of entries in the TrafficScript data.get()/set() storage. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "dataEntries"</li> </ul>

Field Name	Description
data_memory_usage	Number of bytes used in the TrafficScript data.get()/set() storage. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "dataMemoryUsage"</li> </ul>
events_seen	Events seen by the traffic Manager's event handling process. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "eventsSeen"</li> </ul>
hourly_peak_bytes_in_per_second	The peak bytes received from clients per second in the last hour. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "hourlyPeakBytesInPerSecond"</li> </ul>
hourly_peak_bytes_out_per_second	The peak bytes sent to clients per second in the last hour. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "hourlyPeakBytesOutPerSecond"</li> </ul>
hourly_peak_requests_per_second	The peak requests per second in the last hour. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "hourlyPeakRequestsPerSecond"</li> </ul>
hourly_peak_ssl_connections_per_second	The peak ssl connections per second in the last hour. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "hourlyPeakSSLConnectionsPerSecond"</li> </ul>
num_idle_connections	Total number of idle HTTP connections to all nodes (used for future HTTP requests). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numIdleConnections"</li> </ul>
number_child_processes	The number of traffic manager child processes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberChildProcesses"</li> </ul>
number_dnsa_cache_hits	Requests for DNS A records resolved from the traffic manager's local cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberDNSACacheHits"</li> </ul>
number_dnsa_requests	Requests for DNS A records (hostname->IP address) made by the traffic manager. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberDNSARequests"</li> </ul>
number_dnsptr_cache_hits	Requests for DNS PTR records resolved from the traffic manager's local cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberDNSPTRCacheHits"</li> </ul>
number_dnsptr_requests	Requests for DNS PTR records (IP address->hostname) made by the traffic manager. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberDNSPTRRequests"</li> </ul>

Field Name	Description
number_snmp_bad_requests	Malformed SNMP requests received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberSNMPBadRequests"</li> </ul>
number_snmp_get_bulk_requests	SNMP GetBulkRequests received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberSNMPGetBulkRequests"</li> </ul>
number_snmp_get_next_requests	SNMP GetNextRequests received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberSNMPGetNextRequests"</li> </ul>
number_snmp_get_requests	SNMP GetRequests received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberSNMPGetRequests"</li> </ul>
number_snmp_unauthorised_requests	SNMP requests dropped due to access restrictions. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "numberSNMPUnauthorisedRequests"</li> </ul>
ssl_cipher_3des_decrypts	Bytes decrypted with 3DES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipher3DESDecrypts"</li> </ul>
ssl_cipher_3des_encrypts	Bytes encrypted with 3DES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipher3DESEncrypts"</li> </ul>
ssl_cipher_aes_decrypts	Bytes decrypted with AES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherAESDecrypts"</li> </ul>
ssl_cipher_aes_encrypts	Bytes encrypted with AES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherAESEncrypts"</li> </ul>
ssl_cipher_aes_gcm_decrypts	Bytes decrypted with AES-GCM. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherAESGCMDecrypts"</li> </ul>
ssl_cipher_aes_gcm_encrypts	Bytes encrypted with AES-GCM. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherAESGCMEncrypts"</li> </ul>
ssl_cipher_decrypts	Bytes decrypted with a symmetric cipher. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDecrypts"</li> </ul>
ssl_cipher_des_decrypts	Bytes decrypted with DES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDESDecrypts"</li> </ul>

Field Name	Description
ssl_cipher_des_encrypts	Bytes encrypted with DES. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDESEncrypts"</li> </ul>
ssl_cipher_dh_agreements	Number of Diffie Hellman key agreements. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDHAgreements"</li> </ul>
ssl_cipher_dh_generates	Number of Diffie Hellman keys generated. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDHGenerates"</li> </ul>
ssl_cipher_dsa_signs	Number of DSA signing operations. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDSASigns"</li> </ul>
ssl_cipher_dsa_verifies	Number of DSA verifications. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherDSAVerifies"</li> </ul>
ssl_cipher_ecdh_agreements	Number of Elliptic Curve Diffie Hellman key agreements. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherECDHAgreements"</li> </ul>
ssl_cipher_ecdh_generates	Number of Elliptic Curve Diffie Hellman keys generated. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherECDHGenerates"</li> </ul>
ssl_cipher_ecdsa_signs	Number of ECDSA signing operations. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherECDSASigns"</li> </ul>
ssl_cipher_ecdsa_verifies	Number of ECDSA verifications. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherECDSAVerifies"</li> </ul>
ssl_cipher_encrypts	Bytes encrypted with a symmetric cipher. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherEncrypts"</li> </ul>
ssl_cipher_rc4_decrypts	Bytes decrypted with RC4. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRC4Decrypts"</li> </ul>
ssl_cipher_rc4_encrypts	Bytes encrypted with RC4. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRC4Encrypts"</li> </ul>
ssl_cipher_rsa_decrypts	Number of RSA decrypts. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRSADecrypts"</li> </ul>

Field Name	Description
ssl_cipher_rsa_decrypts_external	Number of external RSA decrypts. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRSADecryptsExternal"</li> </ul>
ssl_cipher_rsa_encrypts	Number of RSA encrypts. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRSAEncrypts"</li> </ul>
ssl_cipher_rsa_encrypts_external	Number of external RSA encrypts. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCipherRSAEncryptsExternal"</li> </ul>
ssl_client_cert_expired	Number of times a client certificate has expired. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslClientCertExpired"</li> </ul>
ssl_client_cert_invalid	Number of times a client certificate was invalid. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslClientCertInvalid"</li> </ul>
ssl_client_cert_not_sent	Number of times a client certificate was required but not supplied. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslClientCertNotSent"</li> </ul>
ssl_client_cert_revoked	Number of times a client certificate was revoked. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslClientCertRevoked"</li> </ul>
ssl_connections	Number of SSL connections negotiated. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslConnections"</li> </ul>
ssl_handshake_sslv3	Number of SSLv3 handshakes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslHandshakeSSLv3"</li> </ul>
ssl_handshake_tlsv1	Number of TLSv1.0 handshakes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslHandshakeTLSv1"</li> </ul>
ssl_handshake_tlsv11	Number of TLSv1.1 handshakes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslHandshakeTLSv11"</li> </ul>
ssl_handshake_tlsv12	Number of TLSv1.2 handshakes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslHandshakeTLSv12"</li> </ul>
ssl_handshake_tlsv13	Number of TLSv1.3 handshakes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslHandshakeTLSv13"</li> </ul>

Field Name	Description
ssl_session_id_mem_cache_hit	Number of times the SSL session id was found in the cache and reused. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionIDMemCacheHit"</li> </ul>
ssl_session_id_mem_cache_miss	Number of times the SSL session id was not found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionIDMemCacheMiss"</li> </ul>
sys_cpu_busy_percent	Percentage of time that the CPUs are busy. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysCPUBusyPercent"</li> </ul>
sys_cpu_idle_percent	Percentage of time that the CPUs are idle. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysCPUIdlePercent"</li> </ul>
sys_cpu_system_busy_percent	Percentage of time that the CPUs are busy running system code. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysCPUSystemBusyPercent"</li> </ul>
sys_cpu_user_busy_percent	Percentage of time that the CPUs are busy running user-space code. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysCPUUserBusyPercent"</li> </ul>
sys_fds_free	Number of free file descriptors. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysFDsFree"</li> </ul>
sys_mem_buffered	Buffer memory (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemBuffered"</li> </ul>
sys_mem_free	Free memory (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemFree"</li> </ul>
sys_mem_in_use	Memory used (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemInUse"</li> </ul>
sys_mem_swap_total	Total swap space (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemSwapTotal"</li> </ul>
sys_mem_swapped	Amount of swap space in use (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemSwapped"</li> </ul>
sys_mem_total	Total memory (MBytes). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sysMemTotal"</li> </ul>

Field Name	Description
time_last_config_update	The time (in hundredths of a second) since the configuration of traffic manager was updated (this value will wrap if no configuration changes are made for 497 days). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "timeLastConfigUpdate"</li> </ul>
total_backend_server_errors	Total errors returned from the backend servers. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalBackendServerErrors"</li> </ul>
total_bad_dns_packets	Total number of malformed DNS response packets encountered from the backend servers. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalBadDNSPackets"</li> </ul>
total_bytes_in	Bytes received by the traffic manager from clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "totalBytesIn"</li> </ul>
total_bytes_out	Bytes sent by the traffic manager to clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "totalBytesOut"</li> </ul>
total_conn	Total number of TCP connections received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalConn"</li> </ul>
total_current_conn	Number of TCP connections currently established. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalCurrentConn"</li> </ul>
total_dns_responses	Total number of DNS response packets handled. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalDNSResponses"</li> </ul>
total_requests	Total number of TCP requests received. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalRequests"</li> </ul>
total_transactions	Total number of TCP requests being processed, after applying TPS limits. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "totalTransactions"</li> </ul>
total_udp_bytes_in_dropped	UDP bytes received by the traffic manager from clients but failed to be sent to backend servers. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "totalUDPBytesInDropped"</li> </ul>

Field Name	Description
total_udp_bytes_out_dropped	UDP bytes received by the traffic manager from backend servers but failed to be sent to clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "totalUDPBytesOutDropped"</li> </ul>
up_time	The time (in hundredths of a second) that vTM software has been operational for (this value will wrap if it has been running for more than 497 days). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "upTime"</li> </ul>

### vtm\_traffic\_ips\_ip\_gateway\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/traffic\_ips/ip\_gateway

Ip gateway statistics values.

```
# Example usage of vtm_traffic_ips_ip_gateway_stats data source...
data "vtm_traffic_ips_ip_gateway_stats" "my_traffic_ips_ip_gateway_stats" {}

# Access with...
"${data.vtm_traffic_ips_ip_gateway_stats.my_traffic_ips_ip_gateway_stats.<FIELD_NAME>}"
```

Field Name	Description
arp_message	Number of ARP messages sent for raised Traffic IP Addresses. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPARPMMessage"</li> </ul>
gateway_ping_requests	Number of ping requests sent to the gateway machine. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPGatewayPingRequests"</li> </ul>
gateway_ping_responses	Number of ping responses received from the gateway machine. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPGatewayPingResponses"</li> </ul>
node_ping_requests	Number of ping requests sent to the backend nodes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPNodePingRequests"</li> </ul>
node_ping_responses	Number of ping responses received from the backend nodes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPNodePingResponses"</li> </ul>
number_inet46	The number of traffic IP addresses on this system (includes IPv4 and IPv6 addresses). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPNumberInet46"</li> </ul>



Field Name	Description
number_raised_inet46	The number of traffic IP addresses currently raised on this system (includes IPv4 and IPv6 addresses). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPNumberRaisedInet46"</li> </ul>
ping_response_errors	Number of ping response errors. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPPingResponseErrors"</li> </ul>

### vtm\_cache\_ip\_session\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/ip\_session\_cache

Ip session cache statistics values.

```
# Example usage of vtm_cache_ip_session_cache_stats data source...
data "vtm_cache_ip_session_cache_stats" "my_cache_ip_session_cache_stats" {}

# Access with...
"${data.vtm_cache_ip_session_cache_stats.my_cache_ip_session_cache_stats.<FIELD_NAME>}"
```

Field Name	Description
entries	The total number of IP sessions stored in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheEntries"</li> </ul>
entries_max	The maximum number of IP sessions in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheEntriesMax"</li> </ul>
expiries	Number of times an expired IP session entry has been dropped on lookup. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheExpiries"</li> </ul>
hit_rate	The percentage of IP session lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheHitRate"</li> </ul>
hits	Number of times a IP session entry has been successfully found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheHits"</li> </ul>
lookups	Number of times a IP session entry has been looked up in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheLookups"</li> </ul>

Field Name	Description
misses	Number of times a IP session entry has not been available in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheMisses"</li> </ul>
oldest	The age of the oldest IP session in the cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ipSessionCacheOldest"</li> </ul>

### vtm\_cache\_j2ee\_session\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/j2ee\_session\_cache

J2ee session cache statistics values.

```
# Example usage of vtm_cache_j2ee_session_cache_stats data source...
data "vtm_cache_j2ee_session_cache_stats" "my_cache_j2ee_session_cache_stats" {}

# Access with...
"${data.vtm_cache_j2ee_session_cache_stats.my_cache_j2ee_session_cache_stats.<FIELD_NAME>}"
"
```

Field Name	Description
entries	The total number of J2EE sessions stored in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheEntries"</li> </ul>
entries_max	The maximum number of J2EE sessions in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheEntriesMax"</li> </ul>
expiries	Number of times an expired J2EE session entry has been dropped on lookup. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheExpiries"</li> </ul>
hit_rate	The percentage of J2EE session lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheHitRate"</li> </ul>
hits	Number of times a J2EE session entry has been successfully found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheHits"</li> </ul>
lookups	Number of times a J2EE session entry has been looked up in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheLookups"</li> </ul>

Field Name	Description
misses	Number of times a J2EE session entry has not been available in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheMisses"</li> </ul>
oldest	The age of the oldest J2EE session in the cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "j2eeSessionCacheOldest"</li> </ul>

### vtm\_listen\_ip\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/listen\_ips/\*

Listen ips statistics values.

```
# Example usage of vtm_listen_ip_stats data source...
data "vtm_listen_ip_stats" "my_listen_ip_stats" {
  name = "MyListenIp"
}

# Access with...
"${data.vtm_listen_ip_stats.my_listen_ip_stats.<FIELD_NAME>}"
```

Field Name	Description
bytes_in	Bytes sent to this listening IP. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "listenIPBytesIn"</li> </ul>
bytes_out	Bytes sent from this listening IP. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "listenIPBytesOut"</li> </ul>
current_conn	TCP connections currently established to this listening IP. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "listenIPCurrentConn"</li> </ul>
max_conn	Maximum number of simultaneous TCP connections this listening IP has processed at any one time. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "listenIPMaxConn"</li> </ul>
total_requests	Requests sent to this listening IP. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "listenIPTotalRequests"</li> </ul>

### vtm\_location\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/locations/\*

Locations statistics values.

```
# Example usage of vtm_location_stats data source...
data "vtm_location_stats" "my_location_stats" {
  name = "MyLocation"
}

# Access with...
"${data.vtm_location_stats.my_location_stats.<FIELD_NAME>}"
```

Field Name	Description
load	The mean load metric for this location. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "locationLoad"</li> </ul>
responses	Number of A records that have been altered to point to this location. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "locationResponses"</li> </ul>

### vtm\_network\_interface\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/network\_interface/\*

Network interface statistics values.

```
# Example usage of vtm_network_interface_stats data source...
data "vtm_network_interface_stats" "my_network_interface_stats" {
  name = "MyNetworkInterface"
}

# Access with...
"${data.vtm_network_interface_stats.my_network_interface_stats.<FIELD_NAME>}"
```

Field Name	Description
collisions	The number of collisions reported by this interface. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "interfaceCollisions"</li> </ul>
rx_bytes	Bytes received by this interface. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "interfaceRxBytes"</li> </ul>
rx_errors	The number of receive errors reported by this interface. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "interfaceRxErrors"</li> </ul>
rx_packets	The number of packets received by this interface. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "interfaceRxPackets"</li> </ul>

Field Name	Description
tx_bytes	Bytes transmitted by this interface. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "interfaceTxBytes"</li> </ul>
tx_errors	The number of transmit errors reported by this interface. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "interfaceTxErrors"</li> </ul>
tx_packets	The number of packets transmitted by this interface. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "interfaceTxPackets"</li> </ul>

### vtm\_nodes\_node\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/nodes/node/\*

Node statistics values.

```
# Example usage of vtm_nodes_node_stats data source...
data "vtm_nodes_node_stats" "my_nodes_node_stats" {
  name = "MyNodesNode"
}

# Access with...
"${data.vtm_nodes_node_stats.my_nodes_node_stats.<FIELD_NAME>}"
```

Field Name	Description
bytes_from_node_hi	Bytes received from this node ( high 32bits ). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeBytesFromNodeHi"</li> </ul>
bytes_from_node_lo	Bytes received from this node ( low 32bits ). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeBytesFromNodeLo"</li> </ul>
bytes_to_node_hi	Bytes sent to this node ( high 32bits ). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeBytesToNodeHi"</li> </ul>
bytes_to_node_lo	Bytes sent to this node ( low 32bits ). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeBytesToNodeLo"</li> </ul>
current_conn	Requests currently established to this node. ( does not include idle keepalives ). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeCurrentConn"</li> </ul>

Field Name	Description
current_requests	Connections currently established to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeCurrentRequests"</li> </ul>
errors	Number of timeouts, connection problems and other errors for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeErrors"</li> </ul>
failures	Failures of this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeFailures"</li> </ul>
new_conn	Requests that created a new connection to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeNewConn"</li> </ul>
pooled_conn	Requests that reused an existing pooled/keepalive connection rather than creating a new TCP connection. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodePooledConn"</li> </ul>
port	The port this node listens on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodePort"</li> </ul>
response_max	Maximum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeResponseMax"</li> </ul>
response_mean	Mean response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeResponseMin"</li> </ul>
state	The state of this node. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "nodeState"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"alive": alive(1)</li> <li>"dead": dead(2)</li> <li>"unknown": unknown(3)</li> </ul> </li> </ul>
total_conn	Requests sent to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeTotalConn"</li> </ul>

## vtm\_nodes\_node\_inet46\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/nodes/node\_inet46/\*

Node inet46 statistics values.

```
# Example usage of vtm_nodes_node_inet46_stats data source...
data "vtm_nodes_node_inet46_stats" "my_nodes_node_inet46_stats" {
  name = "MyNodesNodeInet46"
}

# Access with...
"${data.vtm_nodes_node_inet46_stats.my_nodes_node_inet46_stats.<FIELD_NAME>}"
```

Field Name	Description
bytes_from_node	Bytes received from this node. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "nodeInet46BytesFromNode"</li> </ul>
bytes_to_node	Bytes sent to this node. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "nodeInet46BytesToNode"</li> </ul>
current_conn	Current connections established to this node, includes idle connections and connections being established. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46CurrentConn"</li> </ul>
current_requests	Active connections established to this node, does not include idle connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46CurrentRequests"</li> </ul>
errors	Number of timeouts, connection problems and other errors for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46Errors"</li> </ul>
failures	Failures of this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46Failures"</li> </ul>
idle_conns	Number of idle HTTP connections to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46IdleConns"</li> </ul>
new_conn	Requests that created a new connection to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46NewConn"</li> </ul>

Field Name	Description
pooled_conn	Requests that reused an existing pooled/keepalive connection rather than creating a new TCP connection. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46PooledConn"</li> </ul>
port	The port this node listens on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46Port"</li> </ul>
response_max	Maximum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46ResponseMax"</li> </ul>
response_mean	Mean response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46ResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46ResponseMin"</li> </ul>
state	The state of this node. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "nodeInet46State"</li> <li>Permitted values:  "alive": alive(1)  "dead": dead(2)  "unknown": unknown(3) </li> </ul>
total_conn	Requests sent to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "nodeInet46TotalConn"</li> </ul>

### vtm\_per\_location\_service\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/per\_location\_service/\*

Per location service statistics values.

```
# Example usage of vtm_per_location_service_stats data source...
data "vtm_per_location_service_stats" "my_per_location_service_stats" {
  name = "MyPerLocationService"
}

# Access with...
"${data.vtm_per_location_service_stats.my_per_location_service_stats.<FIELD_NAME>}"
```



Field Name	Description
draining	The draining state of this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "perLocationServiceDraining"</li> <li>Permitted values: "draining": draining(1) "active": active(2)</li> </ul>
frontend_state	The frontend state of this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "perLocationServiceFrontendState"</li> <li>Permitted values: "alive": alive(1) "dead": dead(2)</li> </ul>
load	The load metric for this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perLocationServiceLoad"</li> </ul>
monitor_state	The monitor state of this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "perLocationServiceMonitorState"</li> <li>Permitted values: "alive": alive(1) "dead": dead(2)</li> </ul>
responses	Number of A records that have been altered to point to this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perLocationServiceResponses"</li> </ul>
state	The state of this location for this GLB Service. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "perLocationServiceState"</li> <li>Permitted values: "alive": alive(1) "dead": dead(2)</li> </ul>

### vtm\_per\_node\_slm\_per\_node\_service\_level\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/per\_node\_slm/per\_node\_service\_level/\*

Per node service level statistics values.

```
# Example usage of vtm_per_node_slm_per_node_service_level_stats data source...
data "vtm_per_node_slm_per_node_service_level_stats"
"my_per_node_slm_per_node_service_level_stats" {
  name = "MyPerNodeSlmPerNodeServiceLevel"
}
```

```
# Access with...
"${data.vtm_per_node_slm_per_node_service_level_stats.my_per_node_slm_per_node_service_level_stats.<FIELD_NAME>}"
```

Field Name	Description
node_port	The port number of this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelNodePort"</li> </ul>
response_max	Maximum response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelResponseMax"</li> </ul>
response_mean	Mean response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelResponseMin"</li> </ul>
total_conn	Requests handled by this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelTotalConn"</li> </ul>
total_non_conf	Non-conforming requests handled by this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelTotalNonConf"</li> </ul>

### vtm\_per\_node\_slm\_per\_node\_service\_level\_inet46\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/per\_node\_slm/per\_node\_service\_level\_inet46/\*

Per node service level inet46 statistics values.

```
# Example usage of vtm_per_node_slm_per_node_service_level_inet46_stats data source...
data "vtm_per_node_slm_per_node_service_level_inet46_stats"
"my_per_node_slm_per_node_service_level_inet46_stats" {
  name = "MyPerNodeSlmPerNodeServiceLevelInet46"
}
```

```
# Access with...
"${data.vtm_per_node_slm_per_node_service_level_inet46_stats.my_per_node_slm_per_node_service_level_inet46_stats.<FIELD_NAME>}"
```

Field Name	Description
node_port	The port number of this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46NodePort"</li> </ul>
response_max	Maximum response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46ResponseMax"</li> </ul>
response_mean	Mean response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46ResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46ResponseMin"</li> </ul>
total_conn	Requests handled by this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46TotalConn"</li> </ul>
total_non_conf	Non-conforming requests handled by this SLM class to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perNodeServiceLevelInet46TotalNonConf"</li> </ul>

### vtm\_nodes\_per\_pool\_node\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/nodes/per\_pool\_node/\*

Per pool node statistics values.

```
# Example usage of vtm_nodes_per_pool_node_stats data source...
data "vtm_nodes_per_pool_node_stats" "my_nodes_per_pool_node_stats" {
  name = "MyNodesPerPoolNode"
}

# Access with...
"${data.vtm_nodes_per_pool_node_stats.my_nodes_per_pool_node_stats.<FIELD_NAME>}"
```

Field Name	Description
bytes_from_node	Bytes received from this node. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "perPoolNodeBytesFromNode"</li> </ul>
bytes_to_node	Bytes sent to this node. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "perPoolNodeBytesToNode"</li> </ul>
current_conn	Current connections established to a node, includes idle connections and connections being established. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeCurrentConn"</li> </ul>
current_requests	Active connections established to this node, does not include idle connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeCurrentRequests"</li> </ul>
errors	Number of timeouts, connection problems and other errors for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeErrors"</li> </ul>
failures	Failures of this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeFailures"</li> </ul>
idle_conns	Number of idle HTTP connections to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeIdleConns"</li> </ul>
new_conn	Requests that created a new connection to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeNewConn"</li> </ul>
node_port	The port that this node listens on. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeNodePort"</li> </ul>
pooled_conn	Requests that reused an existing pooled/keepalive connection rather than creating a new TCP connection. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodePooledConn"</li> </ul>
response_max	Maximum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeResponseMax"</li> </ul>

Field Name	Description
response_mean	Mean response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeResponseMin"</li> </ul>
state	The state of this node. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "perPoolNodeState"</li> <li>Permitted values:  "alive": alive(1)  "dead": dead(2)  "unknown": unknown(3)  "draining": draining(4)  "drainingtodelete": drainingtodelete(5) </li> </ul>
total_conn	Requests sent to this node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "perPoolNodeTotalConn"</li> </ul>

### vtm\_pool\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/pools/\*

Pools statistics values.

# Example usage of vtm\_pool\_stats data source...

```
data "vtm_pool_stats" "my_pool_stats" {
  name = "MyPool"
}
```

# Access with...

```
"${data.vtm_pool_stats.my_pool_stats.<FIELD_NAME>}"
```

Field Name	Description
algorithm	The load-balancing algorithm the pool uses. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "poolAlgorithm"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"roundrobin": roundrobin(1)</li> <li>"weightedRoundRobin": weightedRoundRobin(2)</li> <li>"perceptive": perceptive(3)</li> <li>"leastConnections": leastConnections(4)</li> <li>"fastestResponseTime": fastestResponseTime(5)</li> <li>"random": random(6)</li> <li>"weightedLeastConnections": weightedLeastConnections(7)</li> </ul> </li> </ul>
bytes_in	Bytes received by this pool from nodes. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolBytesIn"</li> </ul>
bytes_out	Bytes sent by this pool to nodes. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolBytesOut"</li> </ul>
conns_queued	Total connections currently queued to this pool. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolConnsQueued"</li> </ul>
disabled	The number of nodes in this pool that are disabled. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolDisabled"</li> </ul>
draining	The number of nodes in this pool which are draining. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolDraining"</li> </ul>
http1xx_responses	Number of HTTP 1xx responses returned by this pool. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP1xxResponses"</li> </ul>
http2xx_responses	Number of HTTP 2xx responses returned by this pool. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP2xxResponses"</li> </ul>
http3xx_responses	Number of HTTP 3xx responses returned by this pool. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP3xxResponses"</li> </ul>
http4xx_responses	Number of HTTP 4xx responses returned by this pool. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP4xxResponses"</li> </ul>

Field Name	Description
http503_retries	Number of times the pool received an HTTP 503 response from a node and retried it against a different node. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP503Retries"</li> </ul>
http5xx_responses	Number of HTTP 5xx responses returned by this pool. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "poolHTTP5xxResponses"</li> </ul>
max_queue_time	Maximum time a connection was queued for, over the last second. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolMaxQueueTime"</li> </ul>
mean_queue_time	Mean time a connection was queued for, over the last second. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolMeanQueueTime"</li> </ul>
min_queue_time	Minimum time a connection was queued for, over the last second. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolMinQueueTime"</li> </ul>
nodes	The number of nodes registered with this pool. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolNodes"</li> </ul>
persistence	The session persistence method this pool uses <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "poolPersistence"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"none": none(1)</li> <li>"ip": ip(2)</li> <li>"rule": rule(3)</li> <li>"transparent": transparent(4)</li> <li>"applicationCookie": applicationCookie(5)</li> <li>"xZeusBackend": xZeusBackend(6)</li> <li>"ssl": ssl(7)</li> </ul> </li> </ul>
queue_timeouts	Total connections that timed-out while queued. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolQueueTimeouts"</li> </ul>
session_migrated	Sessions migrated to a new node because the desired node was unavailable. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolSessionMigrated"</li> </ul>

Field Name	Description
state	The state of this pool. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "poolState"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"active": active(1)</li> <li>"disabled": disabled(2)</li> <li>"draining": draining(3)</li> <li>"unused": unused(4)</li> <li>"unknown": unknown(5)</li> </ul> </li> </ul>
total_conn	Requests sent to this pool. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "poolTotalConn"</li> </ul>

### vtm\_rule\_authenticator\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/rule\_authenticators/\*

Rule authenticators statistics values.

```
# Example usage of vtm_rule_authenticator_stats data source...
data "vtm_rule_authenticator_stats" "my_rule_authenticator_stats" {
  name = "MyRuleAuthenticator"
}

# Access with...
"${data.vtm_rule_authenticator_stats.my_rule_authenticator_stats.<FIELD_NAME>}"
```

Field Name	Description
errors	Number of connection errors that have occurred when trying to connect to an authentication server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "authenticatorErrors"</li> </ul>
fails	Number of times this Authenticator has failed to authenticate. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "authenticatorFails"</li> </ul>
passes	Number of times this Authenticator has successfully authenticated. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "authenticatorPasses"</li> </ul>
requests	Number of times this Authenticator has been asked to authenticate. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "authenticatorRequests"</li> </ul>



## vtm\_rule\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/rules/\*

Rules statistics values.

```
# Example usage of vtm_rule_stats data source...
data "vtm_rule_stats" "my_rule_stats" {
  name = "MyRule"
}

# Access with...
"${data.vtm_rule_stats.my_rule_stats.<FIELD_NAME>}"
```

Field Name	Description
aborts	Number of times this TrafficScript rule has aborted. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleAborts"</li> </ul>
discards	Number of times this TrafficScript rule has discarded the connection. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleDiscards"</li> </ul>
execution_time_warnings	Number of times this TrafficScript rule has exceeded the execution time warning threshold. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleExecutionTimeWarnings"</li> </ul>
executions	Number of times this TrafficScript rule has been executed. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleExecutions"</li> </ul>
pool_select	Number of times this TrafficScript rule has selected a pool to use. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "rulePoolSelect"</li> </ul>
responds	Number of times this TrafficScript rule has responded directly to the client. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleResponds"</li> </ul>
retries	Number of times this TrafficScript rule has forced the request to be retried. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "ruleRetries"</li> </ul>

## vtm\_service\_level\_monitor\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/service\_level\_monitors/\*

Service level monitors statistics values.

```
# Example usage of vtm_service_level_monitor_stats data source...
data "vtm_service_level_monitor_stats" "my_service_level_monitor_stats" {
  name = "MyServiceLevelMonitor"
}

# Access with...
"${data.vtm_service_level_monitor_stats.my_service_level_monitor_stats.<FIELD_NAME>}"
```

Field Name	Description
conforming	Percentage of requests associated with this SLM class that are conforming <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelConforming"</li> </ul>
current_conns	The number of connections currently associated with this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelCurrentConns"</li> </ul>
is_o_k	Indicates if this SLM class is currently conforming. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "serviceLevelIsOK"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"notok": notok(1)</li> <li>"ok": ok(2)</li> </ul> </li> </ul>
response_max	Maximum response time (ms) in the last second for this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelResponseMax"</li> </ul>
response_mean	Mean response time (ms) in the last second for this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelResponseMean"</li> </ul>
response_min	Minimum response time (ms) in the last second for this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelResponseMin"</li> </ul>
total_conn	Requests handled by this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelTotalConn"</li> </ul>
total_non_conf	Non-conforming requests handled by this SLM class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceLevelTotalNonConf"</li> </ul>

### vtm\_service\_protection\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/service\_protection/\*

Service protection statistics values.

```
# Example usage of vtm_service_protection_stats data source...
data "vtm_service_protection_stats" "my_service_protection_stats" {
  name = "MyServiceProtection"
}

# Access with...
"${data.vtm_service_protection_stats.my_service_protection_stats.<FIELD_NAME>}"
```

Field Name	Description
last_refusal_time	The time (in hundredths of a second) since this service protection class last refused a connection (this value will wrap if no connections are refused in more than 497 days). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtLastRefusalTime"</li> </ul>
refusal_binary	Connections refused by this service protection class because the request contained disallowed binary content. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalBinary"</li> </ul>
refusal_conc10_ip	Connections refused by this service protection class because the top 10 source IP addresses issued too many concurrent connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalConc10IP"</li> </ul>
refusal_conc1_ip	Connections refused by this service protection class because the source IP address issued too many concurrent connections. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalConc1IP"</li> </ul>
refusal_conn_rate	Connections refused by this service protection class because the source IP address issued too many connections within 60 seconds. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalConnRate"</li> </ul>
refusal_ip	Connections refused by this service protection class because the source IP address was banned. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalIP"</li> </ul>
refusal_rfc2396	Connections refused by this service protection class because the HTTP request was not RFC 2396 compliant. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalRFC2396"</li> </ul>

Field Name	Description
refusal_size	Connections refused by this service protection class because the request was larger than the defined limits allowed. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtRefusalSize"</li> </ul>
total_refusal	Connections refused by this service protection class. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "serviceProtTotalRefusal"</li> </ul>

### vtm\_cache\_ssl\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/ssl\_cache

Ssl cache statistics values.

```
# Example usage of vtm_cache_ssl_cache_stats data source...
data "vtm_cache_ssl_cache_stats" "my_cache_ssl_cache_stats" {}

# Access with...
"${data.vtm_cache_ssl_cache_stats.my_cache_ssl_cache_stats.<FIELD_NAME>}"
```

Field Name	Description
entries	The total number of SSL sessions stored in the server cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheEntries"</li> </ul>
entries_max	The maximum number of SSL entries in the server cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheEntriesMax"</li> </ul>
hit_rate	The percentage of SSL server cache lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheHitRate"</li> </ul>
hits	Number of times a SSL entry has been successfully found in the server cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheHits"</li> </ul>
lookups	Number of times a SSL entry has been looked up in the server cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheLookups"</li> </ul>

Field Name	Description
misses	Number of times a SSL entry has not been available in the server cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheMisses"</li> </ul>
oldest	The age of the oldest SSL session in the server cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslCacheOldest"</li> </ul>

### vtm\_ssl\_ocsp\_stapling\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/ssl\_ocsp\_stapling

Ssl ocsp stapling statistics values.

```
# Example usage of vtm_ssl_ocsp_stapling_stats data source...
data "vtm_ssl_ocsp_stapling_stats" "my_ssl_ocsp_stapling_stats" {}

# Access with...
"${data.vtm_ssl_ocsp_stapling_stats.my_ssl_ocsp_stapling_stats.<FIELD_NAME>}"
```

Field Name	Description
cache_count	The number of entries in the OCSP stapling cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingCacheCount"</li> </ul>
counter	The number of outgoing OCSP requests for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingCount"</li> </ul>
failure_count	The number of failed outgoing OCSP requests for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingFailureCount"</li> </ul>
good_count	The number of 'good' OCSP responses for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingGoodCount"</li> </ul>
revoked_count	The number of 'revoked' OCSP responses for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingRevokedCount"</li> </ul>
success_count	The number of successful outgoing OCSP requests for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingSuccessCount"</li> </ul>
unknown_count	The number of 'unknown' OCSP requests for OCSP stapling. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslOcspStaplingUnknownCount"</li> </ul>

## vtm\_cache\_ssl\_session\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/ssl\_session\_cache

Ssl session cache statistics values.

```
# Example usage of vtm_cache_ssl_session_cache_stats data source...
data "vtm_cache_ssl_session_cache_stats" "my_cache_ssl_session_cache_stats" {}

# Access with...
"${data.vtm_cache_ssl_session_cache_stats.my_cache_ssl_session_cache_stats.<FIELD_NAME>}"
```

Field Name	Description
entries	The total number of SSL session persistence entries stored in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheEntries"</li> </ul>
entries_max	The maximum number of SSL session persistence entries in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheEntriesMax"</li> </ul>
hit_rate	The percentage of SSL session persistence lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheHitRate"</li> </ul>
hits	Number of times a SSL session persistence entry has been successfully found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheHits"</li> </ul>
lookups	Number of times a SSL session persistence entry has been looked up in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheLookups"</li> </ul>
misses	Number of times a SSL session persistence entry has not been available in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheMisses"</li> </ul>
oldest	The age of the oldest SSL session in the cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "sslSessionCacheOldest"</li> </ul>

## vtm\_traffic\_ips\_traffic\_ip\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/traffic\_ips/traffic\_ip/\*

Traffic ip statistics values.

```
# Example usage of vtm_traffic_ips_traffic_ip_stats data source...
data "vtm_traffic_ips_traffic_ip_stats" "my_traffic_ips_traffic_ip_stats" {
  name = "MyTrafficIpsTrafficIp"
}

# Access with...
"${data.vtm_traffic_ips_traffic_ip_stats.my_traffic_ips_traffic_ip_stats.<FIELD_NAME>}"
```

Field Name	Description
state	Whether this traffic IP address is currently being hosted by this traffic manager. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "trafficIPState"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"raised": raised(1)</li> <li>"lowered": lowered(2)</li> </ul> </li> </ul>
time	The time (in hundredths of a second) since trafficIPState last changed (this value will wrap if the state hasn't changed for 497 days). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPTime"</li> </ul>

### vtm\_traffic\_ips\_traffic\_ip\_inet46\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/traffic\_ips/traffic\_ip\_inet46/\*

Traffic ip inet46 statistics values.

```
# Example usage of vtm_traffic_ips_traffic_ip_inet46_stats data source...
data "vtm_traffic_ips_traffic_ip_inet46_stats" "my_traffic_ips_traffic_ip_inet46_stats" {
  name = "MyTrafficIpsTrafficIpInet46"
}

# Access with...
"${data.vtm_traffic_ips_traffic_ip_inet46_stats.my_traffic_ips_traffic_ip_inet46_stats.<FIELD_NAME>}"
```

Field Name	Description
state	Whether this traffic IP address is currently being hosted by this traffic manager. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "trafficIPnet46State"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"raised": raised(1)</li> <li>"lowered": lowered(2)</li> </ul> </li> </ul>
time	The time (in hundredths of a second) since trafficIPState last changed (this value will wrap if the state hasn't changed for 497 days). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "trafficIPnet46Time"</li> </ul>

### vtm\_cache\_uni\_session\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/uni\_session\_cache

Uni session cache statistics values.

```
# Example usage of vtm_cache_uni_session_cache_stats data source...
data "vtm_cache_uni_session_cache_stats" "my_cache_uni_session_cache_stats" {}

# Access with...
"${data.vtm_cache_uni_session_cache_stats.my_cache_uni_session_cache_stats.<FIELD_NAME>}"
```

Field Name	Description
entries	The total number of universal sessions stored in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheEntries"</li> </ul>
entries_max	The maximum number of universal sessions in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheEntriesMax"</li> </ul>
expiries	Number of times an expired universal session entry has been dropped on lookup. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheExpiries"</li> </ul>
hit_rate	The percentage of universal session lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheHitRate"</li> </ul>
hits	Number of times a universal session entry has been successfully found in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheHits"</li> </ul>



Field Name	Description
lookups	Number of times a universal session entry has been looked up in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheLookups"</li> </ul>
misses	Number of times a universal session entry has not been available in the cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheMisses"</li> </ul>
oldest	The age of the oldest universal session in the cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "uniSessionCacheOldest"</li> </ul>

### vtm\_extras\_user\_counters\_32\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/extras/user\_counters\_32

User counters 32 statistics values.

```
# Example usage of vtm_extras_user_counters_32_stats data source...
data "vtm_extras_user_counters_32_stats" "my_extras_user_counters_32_stats" {}

# Access with...
"${data.vtm_extras_user_counters_32_stats.my_extras_user_counters_32_stats.<FIELD_NAME>}"
```

Field Name	Description
counter	The value of the user counter. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "userCounterValue"</li> </ul>

### vtm\_extras\_user\_counters\_64\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/extras/user\_counters\_64

User counters 64 statistics values.

```
# Example usage of vtm_extras_user_counters_64_stats data source...
data "vtm_extras_user_counters_64_stats" "my_extras_user_counters_64_stats" {}

# Access with...
"${data.vtm_extras_user_counters_64_stats.my_extras_user_counters_64_stats.<FIELD_NAME>}"
```

Field Name	Description
counter	The value of the 64-bit user counter. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "userCounter64Value"</li> </ul>

### vtm\_virtual\_server\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/virtual\_servers/\*

Virtual servers statistics values.

```
# Example usage of vtm_virtual_server_stats data source...
data "vtm_virtual_server_stats" "my_virtual_server_stats" {
  name = "MyVirtualServer"
}

# Access with...
"${data.vtm_virtual_server_stats.my_virtual_server_stats.<FIELD_NAME>}"
```

Field Name	Description
auth_saml_redirects	Number of times a user agent was redirected to SAML Identity Provider. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSamlRedirects"</li> </ul>
auth_saml_responses	Number of times a SAML Response was processed. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSamlResponses"</li> </ul>
auth_saml_responses_accepted	Number of times a SAML Response was accepted. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSamlResponsesAccepted"</li> </ul>
auth_saml_responses_rejected	Number of times a SAML Response was rejected. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSamlResponsesRejected"</li> </ul>
auth_sessions_created	Number of times an authentication session was created. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSessionsCreated"</li> </ul>
auth_sessions_rejected	Number of times an authentication session was rejected. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSessionsRejected"</li> </ul>
auth_sessions_used	Number of times an authentication session was used. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverAuthSessionsUsed"</li> </ul>

Field Name	Description
bytes_in	Bytes received by this virtual server from clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverBytesIn"</li> </ul>
bytes_out	Bytes sent by this virtual server to clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverBytesOut"</li> </ul>
cert_status_requests	Number of incoming TLS handshakes for this virtual server with certificate status requests. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverCertStatusRequests"</li> </ul>
cert_status_responses	Number of incoming TLS handshakes for this virtual server to which certificate status responses were attached. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverCertStatusResponses"</li> </ul>
connect_timed_out	Connections closed by this virtual server because the 'connect_timeout' interval was exceeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverConnectTimedOut"</li> </ul>
connection_errors	Number of transaction or protocol errors in this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverConnectionErrors"</li> </ul>
connection_failures	Number of connection failures in this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverConnectionFailures"</li> </ul>
current_conn	TCP connections currently established to this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverCurrentConn"</li> </ul>
data_timed_out	Connections closed by this virtual server because the 'timeout' interval was exceeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverDataTimedOut"</li> </ul>
direct_replies	Direct replies from this virtual server, without forwarding to a node. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverDirectReplies"</li> </ul>
discard	Connections discarded by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverDiscard"</li> </ul>
gzip	Responses which have been compressed by content compression. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverGzip"</li> </ul>

Field Name	Description
gzip_bytes_saved	Bytes of network traffic saved by content compression. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverGzipBytesSaved"</li> </ul>
http1xx_responses	Number of HTTP 1xx responses returned by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTP1xxResponses"</li> </ul>
http2xx_responses	Number of HTTP 2xx responses returned by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTP2xxResponses"</li> </ul>
http3xx_responses	Number of HTTP 3xx responses returned by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTP3xxResponses"</li> </ul>
http4xx_responses	Number of HTTP 4xx responses returned by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTP4xxResponses"</li> </ul>
http5xx_responses	Number of HTTP 5xx responses returned by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTP5xxResponses"</li> </ul>
http_cache2xx_responses	Number of HTTP 2xx responses returned from webcache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPCache2xxResponses"</li> </ul>
http_cache3xx_responses	Number of HTTP 3xx responses returned from webcache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPCache3xxResponses"</li> </ul>
http_cache4xx_responses	Number of HTTP 4xx responses returned from webcache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPCache4xxResponses"</li> </ul>
http_cache5xx_responses	Number of HTTP 5xx responses returned from webcache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPCache5xxResponses"</li> </ul>
http_cache_hit_rate	Percentage hit rate of the web cache for this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverHttpCacheHitRate"</li> </ul>
http_cache_hits	HTTP responses sent directly from the web cache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverHttpCacheHits"</li> </ul>

Field Name	Description
http_cache_lookups	HTTP requests that are looked up in the web cache by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverHttpCacheLookups"</li> </ul>
http_generated2xx_responses	Number of HTTP 2xx responses generated by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPGenerated2xxResponses"</li> </ul>
http_generated3xx_responses	Number of HTTP 3xx responses generated by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPGenerated3xxResponses"</li> </ul>
http_generated4xx_responses	Number of HTTP 4xx responses generated by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPGenerated4xxResponses"</li> </ul>
http_generated5xx_responses	Number of HTTP 5xx responses generated by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPGenerated5xxResponses"</li> </ul>
http_rewrite_cookie	HTTP Set-Cookie headers, supplied by a node, that have been rewritten. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverHttpRewriteCookie"</li> </ul>
http_rewrite_location	HTTP Location headers, supplied by a node, that have been rewritten. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverHttpRewriteLocation"</li> </ul>
http_server1xx_responses	Number of HTTP 1xx responses returned from a backend server or TrafficScript rule by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPServer1xxResponses"</li> </ul>
http_server2xx_responses	Number of HTTP 2xx responses returned from a backend server or TrafficScript rule by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPServer2xxResponses"</li> </ul>
http_server3xx_responses	Number of HTTP 3xx responses returned from a backend server or TrafficScript rule by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPServer3xxResponses"</li> </ul>
http_server4xx_responses	Number of HTTP 4xx responses returned from a backend server or TrafficScript rule by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPServer4xxResponses"</li> </ul>

Field Name	Description
http_server5xx_responses	<p>Number of HTTP 5xx responses returned from a backend server or TrafficScript rule by this virtual server.</p> <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverHTTPServer5xxResponses"</li> </ul>
keepalive_timed_out	<p>Connections closed by this virtual server because the 'keepalive_timeout' interval was exceeded.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverKeepaliveTimedOut"</li> </ul>
max_conn	<p>Maximum number of simultaneous TCP connections this virtual server has processed at any one time.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverMaxConn"</li> </ul>
max_duration_timed_out	<p>Connections closed by this virtual server because the 'max_transaction_duration' interval was exceeded.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverMaxDurationTimedOut"</li> </ul>
port	<p>The port the virtual server listens on.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverPort"</li> </ul>
processing_timed_out	<p>Connections closed by this virtual server because the 'timeout' interval was exceeded while waiting for rules or external processing.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverProcessingTimedOut"</li> </ul>

Field Name	Description
protocol	<p>The protocol the virtual server is operating.</p> <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>SNMP name: "virtualserverProtocol"</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"http": http(1)</li> <li>"https": https(2)</li> <li>"ftp": ftp(3)</li> <li>"imaps": imaps(4)</li> <li>"imapv2": imapv2(5)</li> <li>"imapv3": imapv3(6)</li> <li>"imapv4": imapv4(7)</li> <li>"pop3": pop3(8)</li> <li>"pop3s": pop3s(9)</li> <li>"smtp": smtp(10)</li> <li>"ldap": ldap(11)</li> <li>"ldaps": ldaps(12)</li> <li>"telnet": telnet(13)</li> <li>"sslforwarding": sslforwarding(14)</li> <li>"udpstreaming": udpstreaming(15)</li> <li>"udp": udp(16)</li> <li>"dns": dns(17)</li> <li>"genericserverfirst": genericserverfirst(18)</li> <li>"genericclientfirst": genericclientfirst(19)</li> <li>"dnstcp": dnstcp(20)</li> <li>"sipudp": sipudp(21)</li> <li>"siptcp": siptcp(22)</li> <li>"rtsp": rtsp(23)</li> <li>"stream": stream(24)</li> <li>"l4acceltcp": l4acceltcp(25)</li> <li>"l4acceludp": l4acceludp(26)</li> <li>"l4accelgeneric": l4accelgeneric(27)</li> <li>"l4accelstateless": l4accelstateless(28)</li> </ul> </li> </ul>
sip_rejected_requests	<p>Number of SIP requests rejected due to them exceeding the maximum amount of memory allocated to the connection.</p> <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverSIPRejectedRequests"</li> </ul>

Field Name	Description
sip_total_calls	Total number of SIP INVITE requests seen by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverSIPTotalCalls"</li> </ul>
ssl_cache_lookup	Number of times a lookup for an existing SSL session was performed. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslCacheLookup"</li> </ul>
ssl_cache_miss	Number of times a lookup failed to find an existing SSL session. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslCacheMiss"</li> </ul>
ssl_cache_rejected	Number of times an SSL session was found in the cache but rejected and not resumed. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslCacheRejected"</li> </ul>
ssl_cache_resumed	Number of times an SSL session was resumed from the cache. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslCacheResumed"</li> </ul>
ssl_cache_saved	Number of times an SSL session was saved to the cache. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslCacheSaved"</li> </ul>
ssl_hello_retry_requested	Number of times a HelloRetryRequest message was sent to TLS clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslHelloRetryRequested"</li> </ul>
ssl_new_session	Number of times a new SSL session was created. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslNewSession"</li> </ul>
ssl_ticket_expired	Number of SSL session tickets that were rejected because they had expired. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketExpired"</li> </ul>
ssl_ticket_issued	Number of SSL session tickets that were issued to clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketIssued"</li> </ul>
ssl_ticket_key_not_found	Number of SSL session tickets that could not be decrypted because the ticket key they referenced could not be found. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketKeyNotFound"</li> </ul>
ssl_ticket_received	Number of SSL session tickets received. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketReceived"</li> </ul>



Field Name	Description
ssl_ticket_rejected	Number of SSL session tickets that were rejected for a reason other than because they had expired. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketRejected"</li> </ul>
ssl_ticket_resumed	Number of SSL session tickets that were successfully used to resume a session. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverSslTicketResumed"</li> </ul>
total_dgram	UDP datagrams processed by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverTotalDgram"</li> </ul>
total_http1_requests	HTTP/1.x Requests received by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverTotalHTTP1Requests"</li> </ul>
total_http2_requests	HTTP/2 Requests received by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverTotalHTTP2Requests"</li> </ul>
total_http_requests	HTTP Requests received by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverTotalHTTPRequests"</li> </ul>
total_requests	Requests received by this virtual server. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverTotalRequests"</li> </ul>
udp_bytes_in_dropped	UDP bytes received by this virtual server from clients but failed to be sent to backend servers. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverUDPBytesInDropped"</li> </ul>
udp_bytes_out_dropped	UDP bytes received by this virtual server from backend servers but failed to be sent to clients. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "virtualserverUDPBytesOutDropped"</li> </ul>
udp_timed_out	Connections closed by this virtual server because the 'udp_timeout' interval was exceeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "virtualserverUdpTimedOut"</li> </ul>

### vtm\_cache\_web\_cache\_stats

URI Endpoint: /api/tm/8.1/status/local\_tm/statistics/cache/web\_cache

Web cache statistics values.

```
# Example usage of vtm_cache_web_cache_stats data source...
data "vtm_cache_web_cache_stats" "my_cache_web_cache_stats" {}

# Access with...
"${data.vtm_cache_web_cache_stats.my_cache_web_cache_stats.<FIELD_NAME>}"
```

Field Name	Description
entries	The number of items in the web cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheEntries"</li> </ul>
hit_rate	The percentage of web cache lookups that succeeded. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheHitRate"</li> </ul>
hits	Number of times a page has been successfully found in the web cache. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheHits"</li> </ul>
lookups	Number of times a page has been looked up in the web cache. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheLookups"</li> </ul>
max_entries	The maximum number of items in the web cache. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheMaxEntries"</li> </ul>
mem_maximum	The maximum amount of memory the web cache can use in kilobytes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheMemMaximum"</li> </ul>
mem_used	Total memory used by the web cache in kilobytes. <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheMemUsed"</li> </ul>
misses	Number of times a page has not been found in the web cache. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheMisses"</li> </ul>
oldest	The age of the oldest item in the web cache (in seconds). <ul style="list-style-type: none"> <li>Type: UInt</li> <li>SNMP name: "webCacheOldest"</li> </ul>
url_store_allocated	Amount of allocated space in the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreAllocated"</li> </ul>
url_store_free	Amount of free space in the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreFree"</li> </ul>

Field Name	Description
url_store_size	Total amount of space in the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreSize"</li> </ul>
url_store_total_allocations	Total number of allocations for the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreTotalAllocations"</li> </ul>
url_store_total_failures	Total number of allocation failures for the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreTotalFailures"</li> </ul>
url_store_total_frees	Total number of blocks freed in the web cache URL store. <ul style="list-style-type: none"> <li>Type: UInt64</li> <li>SNMP name: "webCacheURLStoreTotalFrees"</li> </ul>

## System Information

System information data sources allow Terraform to access information about the current state of a Traffic Manager, including errors and service statuses.

```
data "vtm_state" "state_information" {}
```

### vtm\_backups\_full

URI Endpoint: /api/tm/8.1/status/local\_tm/backups/full/\*

Full backups.

```
# Example usage of vtm_backups_full data source...
```

```
data "vtm_backups_full" "my_backups_full" {
  name = "MyBackupsFull"
}
```

```
# Access with...
```

```
"${data.vtm_backups_full.my_backups_full.<FIELD_NAME>}"
```

Property	Description
name	Name of the object <ul style="list-style-type: none"> <li>Type: string</li> <li>Required: true</li> </ul>
backup_description	Description of the backup <ul style="list-style-type: none"> <li>Type: String</li> </ul>

Property	Description
backup_time_stamp	Time the backup was created. Expressed as a UTC value. <ul style="list-style-type: none"> <li>Type: Int</li> </ul>
backup_version	Version of the traffic manager used to create the backup <ul style="list-style-type: none"> <li>Type: String</li> </ul>

### vtm\_information

URI Endpoint: /api/tm/8.1/status/local\_tm/information

Static information for the system.

```
# Example usage of vtm_information data source...
data "vtm_information" "my_information" {}

# Access with...
"${data.vtm_information.my_information.<FIELD_NAME>}"
```

Property	Description
information_platform	The type of platform on which the Traffic Manager instance is running on. <ul style="list-style-type: none"> <li>Type: String</li> </ul>
information_tm_version	Version number of the Traffic Manager instance. <ul style="list-style-type: none"> <li>Type: String</li> </ul>
information_uuid	The universally unique identifier for the Traffic Manager instance. <ul style="list-style-type: none"> <li>Type: String</li> </ul>

### vtm\_state

URI Endpoint: /api/tm/8.1/status/local\_tm/state

State information for the traffic manager.

```
# Example usage of vtm_state data source...
data "vtm_state" "my_state" {}

# Access with...
"${data.vtm_state.my_state.<FIELD_NAME>}"
```

Property	Description
state_error_level	The error_level of the traffic manager. <ul style="list-style-type: none"> <li>Type: Enum(String)</li> <li>Permitted values: <ul style="list-style-type: none"> <li>"ok": System has no problems</li> <li>"warn": System has minor issues</li> <li>"error": System has major issues</li> <li>"fatal": System has issues which causes it to die/crash/fail to startup</li> </ul> </li> </ul>
state_errors	List of configuration errors for the traffic manager <ul style="list-style-type: none"> <li>Type: Set(String)</li> </ul>
state_failed_nodes	A table of nodes which have failed on the traffic manager <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>node (String): A node which has failed (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>pools (Set(String)): List of pools which use this node.</li> </ul> </li> </ul>
state_failed_nodes_json	A Traffic Manager REST-compatible JSON representation of the "state_failed_nodes" table property. Use this field with the "vtm_state_failed_nodes_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
state_license	Current active license or Community_Edition <ul style="list-style-type: none"> <li>Type: String</li> </ul>
state_pools	<ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): Name of the pool (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>failure_pool (String): Failure pool associated with this pool.</li> <li>active_nodes (Set(String)): List of nodes which are in the active state.</li> <li>draining_nodes (Set(String)): List of nodes which are in the draining state.</li> <li>disabled_nodes (Set(String)): List of nodes which are in the disabled state.</li> </ul> </li> </ul>

Property	Description
state_pools_json	A Traffic Manager REST-compatible JSON representation of the "state_pools" table property. Use this field with the "vtm_state_pools_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>
state_tip_errors	List of traffic IP errors for the traffic manager <ul style="list-style-type: none"> <li>Type: Set(String)</li> </ul>
state_virtual_servers	A table of virtual server status <ul style="list-style-type: none"> <li>Type: Table</li> <li>Required: false</li> <li>Primary key: <ul style="list-style-type: none"> <li>name (String): Name of the virtual server (Required)</li> </ul> </li> <li>Sub keys: <ul style="list-style-type: none"> <li>pool (String): Pool associated with this virtual server.</li> <li>port (UInt): Port the virtual server listens on.</li> <li>throughput (UInt64): Through put for the virtual server.</li> <li>ts_redirect_pools (Set(String)): List of pools which may be selected by Traffic Script Request Rules.</li> </ul> </li> </ul>
state_virtual_servers_json	A Traffic Manager REST-compatible JSON representation of the "state_virtual_servers" table property. Use this field with the "vtm_state_virtual_servers_table" data source for dynamic table generation from input variables. <ul style="list-style-type: none"> <li>Type: String</li> <li>Required: false</li> <li>Default value: &lt;none&gt;</li> </ul>

## Table Field Data Sources

When hard-coding a Terraform template, you can include nested data structures that represent the table structures used by some Traffic Manager configuration fields. For example, the "nodes\_table" field of the "vtm\_pool" resource can have multiple rows specified as shown in the code sample that follows:

```
resource "vtm_pool" "my_pool" {
  name = "MyPool"
  nodes_table {
    node = "192.0.2.1:80"
    state = "active"
  }
  nodes_table {
    node = "192.0.2.2:80"
    state = "active"
  }
  nodes_table {
    node = "192.0.2.3:80"
    state = "disabled"
  }
}
```

```
}
```

To add a new row to the "nodes\_table" field through this method, add another "nodes\_table {}" block with the relevant field values. Often, however, it is desirable to assign configuration data to a template using input variables, although currently there is no supported method for passing in nested structures.

To facilitate the use-case of passing table configuration into a template through input variables, the Traffic Manager provides table field data sources that can accept standard list variables as inputs to each column of the table field. To identify the corresponding table property, each table field data source uses the naming convention "<RESOURCE\_NAME>\_<TABLE\_FIELD\_NAME>\_table". Each data source supports the same fields, with the same default values, as the corresponding table and returns a JSON string that can be supplied to a configuration resource. Each data source also contains a "json" field that automatically computes the properly-formatted JSON representation of the table row. These json values can then be assembled into a JSON list and assigned to a special field on the resource object called "<TABLE\_FIELD\_NAME>\_json".

For example, using the "vtm\_pool" -> "nodes\_table" data from the preceding code sample, use the table field data source "vtm\_pool\_nodes\_table\_table" to apply the same configuration through input variables:

```
# In the terraform.tfvars file (or equivalent through the command line)
node_list = ["192.0.2.1:80", "192.0.2.2:80", "192.0.2.3:80"]
node_state_list = ["active", "active", "disabled"]

# In the Terraform template...
variable "node_list" {
  type = list
}

variable "node_state_list" {
  type = list
}

data "vtm_pool_nodes_table_table" "nodes_table" {
  count = "${length(var.node_list)}"
  node = "${var.node_list[count.index]}"
  state = "${var.node_state_list[count.index]}"
}

resource "vtm_pool" "my_pool" {
  name = "MyPool"
  nodes_table_json = "${join(",", data.vtm_pool_nodes_table_table.nodes_table.*.json)}"
}
```

The following list contains all available table field data sources:

- vtm\_action\_arguments\_table
- vtm\_appliance\_nat\_many\_to\_one\_all\_ports\_table
- vtm\_appliance\_nat\_many\_to\_one\_port\_locked\_table
- vtm\_appliance\_nat\_one\_to\_one\_table
- vtm\_appliance\_nat\_port\_mapping\_table

- vtm\_custom\_string\_lists\_table
- vtm\_glb\_service\_dnssec\_keys\_table
- vtm\_glb\_service\_location\_settings\_table
- vtm\_global\_settings\_appliance\_returnpath\_table
- vtm\_log\_export\_metadata\_table
- vtm\_monitor\_arguments\_table
- vtm\_pool\_nodes\_table\_table
- vtm\_state\_failed\_nodes\_table
- vtm\_state\_pools\_table
- vtm\_state\_virtual\_servers\_table
- vtm\_traffic\_ip\_group\_ip\_mapping\_table
- vtm\_traffic\_manager\_appliance\_card\_table
- vtm\_traffic\_manager\_appliance\_sysctl\_table
- vtm\_traffic\_manager\_hosts\_table
- vtm\_traffic\_manager\_if\_table
- vtm\_traffic\_manager\_ip\_table
- vtm\_traffic\_manager\_routes\_table
- vtm\_traffic\_manager\_trafficip\_table
- vtm\_user\_group\_permissions\_table
- vtm\_virtual\_server\_ocsp\_issuers\_table
- vtm\_virtual\_server\_profile\_table
- vtm\_virtual\_server\_server\_cert\_host\_mapping\_table



# Example Deployment Template

---

This chapter contains an example template to demonstrate deployment of your Traffic Manager infrastructure through the Terraform provider.

```

provider "vtm" {
  base_url = "https://192.0.2.100:9070/api"
  username = "admin"
  password = "admin"
  verify_ssl_cert = false
}

#####
# TIP GROUPS
#####

resource "vtm_traffic_ip_group" "web_tip" {
  name = "web_tip"
  enabled = true
  ipaddresses = ["192.0.2.200"]
  machines = ["192.0.2.250"]
}

#####
# POOLS
#####

resource "vtm_pool" "html_pool" {
  name = "html_pool"
  load_balancing_algorithm = "weighted_least_connections"
  nodes_table {
    node = "192.0.2.10:80"
    weight = 1
  }
  nodes_table {
    node = "192.0.2.11:80"
    weight = 3
  }
  nodes_table {
    node = "192.0.2.12:80"
    weight = 3
    state = "disabled"
  }
}

resource "vtm_pool" "images_pool" {
  name = "images_pool"
  nodes_table {
    node = "192.0.2.20:80"
  }
}

```

```

nodes_table {
  node = "192.0.2.21:80"
}
}

#####
# RULES
#####

resource "vtm_rule" "use_images_pool" {
  name = "UseImagesPool"
  content = <<EOF
if(http.getPath() == "/images") {
  pool.use("${vtm_pool.images_pool.name}");
}
EOF
}

resource "vtm_rule" "redirect_to_https" {
  name = "RedirectToHTTPS"
  content = <<EOF
$hostHeader = http.getHostHeader();
http.redirect("https://" . $hostHeader);
EOF
}

#####
# SSL CERTS
#####

resource "vtm_server_key" "ssl_cert" {
  name      = "SSL-Cert"
  note      = ""
  private   = <<EOF
-----BEGIN RSA PRIVATE KEY-----
MIIEogIBAAKCAQEApEPRlK+xJbQfUenl9H4nLEkQaH5L8/9F+pcjJW14EdlSkI3s
6bPe+eGtWf0XSzDXzOqAufERWrKyhw21c+UYoTA64i43T9nwvlRtxXxcok+VqmQC
HMCT5V7d82DKXnEEE6J5LwmHo48MsaQsBrjeyGVA8n40JDoM3qC1IlyxEoqouRf8
5eYWYAUGL1PSnMT71fZy6VUAaeFRVRwREc3RFxkKa9GEraCaDGp7jfmnrNH1A8Pn
ns3um6kWuRHIliewFhUmclgTxoQurVTtQh/FQViA6UbQtVCcScYarjqo9dgcuz7R
ABCEBnTx7gcmDeB6VR7luN3MGTc8gupJOjNEawIDAQABAoIBAH1+s8/ulu7VJ+k1
P+EzQQZQOwnUXHORuZfrvuI4hrpK1E91ZK+a7JGvceJSjjowNpz+oHufotrZHv6YG
bLQ4uZvXCWZrWnvULa0I01wjHHzssj0mo+/SPFGFdlJhv6xUmPhQzqMMwGcoEfJ+
BBZAvH6p7Xyt/bDhws0TPoYUDB2yg0VN0lgdjg3S/xR2od8ggZBWOW5DV3S0j5ta
Q1LojcDgHh6MQ1VOHIUHhllHUKuhlHLhukUHUKhkHKLldMPUrghmJCaGVSDGr9+
2pjPhz04fp7CZ264rmUyHCGhwhjz4FH1KqkVM5LxoKhbOobwivhrhI/g4skgPxfb
Y7lwxsECgYEA2JgEzk6tnCZO8Ufir+FSvCLWmFJAlHuFasUy8B0jPZiLRNxt88iz
BhSHIJadC09fmMF6WdE1KTZgNnMnm675FDSSKjk12DhGzU/zRqaBh0NE/Lemqe8/
wnw3oLd3Z++zcrJj7P6NCoQN8BEIJ0voiqje+ppEkumftQ8VEHxGEuECgYEAwian
F4W5Z0Y2EEnkwbz1w+Fzaur0g3balipfjt/Up4xWDnsMdp7xPzny2rGzfqfOpUOF
dRR2YwXGhActMpo+Nblg9/0YpJ500RDxMMYEj8bYjTaEFJQ4a7SSvki35pFM0/pT
smhyWA04U/ttbs3+XosN335JsbcopZft3zF3zMsCgYBWDniCXAjgs1vURAJlGuKb
e6AV30BnfnhxBq8Jdhpus5V5Obe6D661HVIEobL+BmhuMhlhzFy55i/jeq9ua398

```

```

p7pn/x998Yb0FlsLbCa0zoZ/fpyKklOcM76eraaUtklumKb5R95UGknLY4kAzAk1
5ojJuzeZw5cWr/JnnWjeIQKBgE09explmBpPI4kdbMXrADeaxQk/SmHW8iWV3AiC
Yh76RO5j49PT7XSDAGwjEE8OQbccAsdOib7heFXkXq3eEWvcQYjHh3tOkxjtzZbi
4MO2j0a27pslUMEAYPstB4TSP6eByrSKuxruv38h4yqXB2Djn3RP0M/EF4axvZfp
HUK7AoGAL2OaOtwy5lk/Oc6bwuPjTAR0wmBX9zelgsLIPiON2jHY287syVzh7lPm
vQGRYZlbgMseXj++s9nQJ24gLTokX2FwGioKvXwFX1ujah7ccJR9iVEwKpQMtDY5
cTiUZkme5oO3Idw6IO115A2EB1/BPpoWBP0M+y2BQYuTGk8F2LU=
-----END RSA PRIVATE KEY-----
EOF
  public = <<EOF
-----BEGIN CERTIFICATE-----
MIIDIDCCAgigAwIBAgIJAMD7f7Ux921VMA0GCSqGSIb3DQEBCwUAMD4xCzAJBgNV
BAYTAkdCMRAwDgYDVQQHEwdkYXNkYXNkMQ0wCwYDVQQKEwRCbGFoMQ4wDAYDVQQD
EwVhLmIuYzAeFw0xODAxMjUyMzI3MDNaFw0yODAxMjUyMzI3MDNaMD4xCzAJBgNV
BAYTAkdCMRAwDgYDVQQHEwdkYXNkYXNkMQ0wCwYDVQQKEwRCbGFoMQ4wDAYDVQQD
EwVhLmIuYzCCASIwDQYJKoZIhvcNAQEHkdakdhJLUihiQoCggEBAKRD0ZSvsSW0
H1Hp5fR+JyxJEWERWEewfrRtIyVteBHZUpCN7Omz3vnhrVn9F0sw18zqgLnxEVqy
socNtXPlGKEwOuIuN0/Z8L5UbcV8XKJPlapqghzAk+Ve3fNgyl5xBBoies8Jh6OP
DLGkLAA43shlQPJ+NCQ6DN6gtSJWMRKKqLkX/OXmFmALoC9T0pze+9X2culVAGnh
UVUcERHN0RcZCmvRhK2gmgxqe435nazR9QPD557N7pupFrkRyJYnsBYVJnNYE8aE
Lq1U7UIfxUFYgOlG0LVQnEnGGq46qPXYHLS+0QAQhAZ08e4HJg3gelUe5bjdzBk3
PILqSTozRGsCAwEAAAMhMB8wHQYDVR0OBBYEFOSQUOxzga482TRQfgcvWsOXHu3k
MA0GCSqGSIb3DQEBCwUAA4IBAQAoufUIugke4ZHRAXYmgM5cUX1MbBU5S71u+Ao
79RfGkDL1kfVpdAoQx1/EoWc7LRIzvbuIZu6BiarU+/Te6mirmjF+dFdCfEka7cY
ZR5/BvU/+xJNEFDz2bEL0f4LTKnEiloEcUshAt3vaqRdBGNt3vvpJ5FjyaDXjmpA
idvAjKqXEbUUBgt0kWuaQU8CDCv5FiGr9XhmK8YnoABCsyALbF+NP41EyUfZzt0Z
bj25+V9mexgCGRR6HJI9whhz33v51SXjxlAX5vsDiXRhfLST7MBGame6nqew2k9
cMbhPHfTIYYM1ijaqFU/LEXOQ6jTieldVIvC0KVSue7+eQtn
-----END CERTIFICATE-----
EOF
  request = ""
}

#####
# VIRTUAL SERVERS
#####

resource "vtm_virtual_server" "http_vs" {
  name = "HTTP"
  enabled = true
  pool = "discard"
  port = 80
  listen_on_any = false
  listen_on_traffic_ips = ["${vtm_traffic_ip_group.web_tip.name}"]
  request_rules = [
    "${vtm_rule.use_images_pool.name}",
    "${vtm_rule.redirect_to_https.name}",
  ]
  web_cache_enabled = true
  web_cache_refresh_time = 30
}

resource "vtm_virtual_server" "https_vs" {
  name = "HTTPS"

```

```
enabled = true
pool = "${vtm_pool.html_pool.name}"
port = 443
listen_on_any = false
listen_on_traffic_ips = ["${vtm_traffic_ip_group.web_tip.name}"]
ssl_decrypt = true
ssl_server_cert_default = "${vtm_server_key.ssl_cert.name}"
}
```