



**VERITAS™**

# Veritas InfoScale™ Chef Deployment Guide

Last updated: 2018-09-19

## Table of Contents

INTRODUCTION.....	3
CHEF ENVIRONMENT.....	3
USING CHEF TO INSTALL AND CONFIGURE INFOSCALE IN A LINUX, SOLARIS, OR AIX ENVIRONMENT.....	4
<i>Prerequisites</i> .....	5
<i>Step 1: Downloading the Veritas InfoScale cookbook from the supermarket</i> .....	5
<i>Step 2: Specifying operations in a recipe file</i> .....	6
<i>Step 3: Creating roles and defining attributes</i> .....	7
<i>Step 4: Running the Chef-client on the Chef node</i> .....	40
EXAMPLE OF ATTRIBUTES FOR INSTALLATION AND CONFIGURATION.....	40

## Introduction

You can now deploy and configure the Veritas InfoScale product suite in your environment using Chef. Chef's powerful automation platform transforms infrastructure into code and allows you to automate various deployment and configuration operations performed in Veritas InfoScale.

### Supported Chef Version

The procedure in this guide is supported with *Chef Client 12.5.1* and later versions.

### Supported platforms

You can use the *veritas\_infoscale* cookbook to perform operations on Veritas InfoScale 7.0 and later versions.

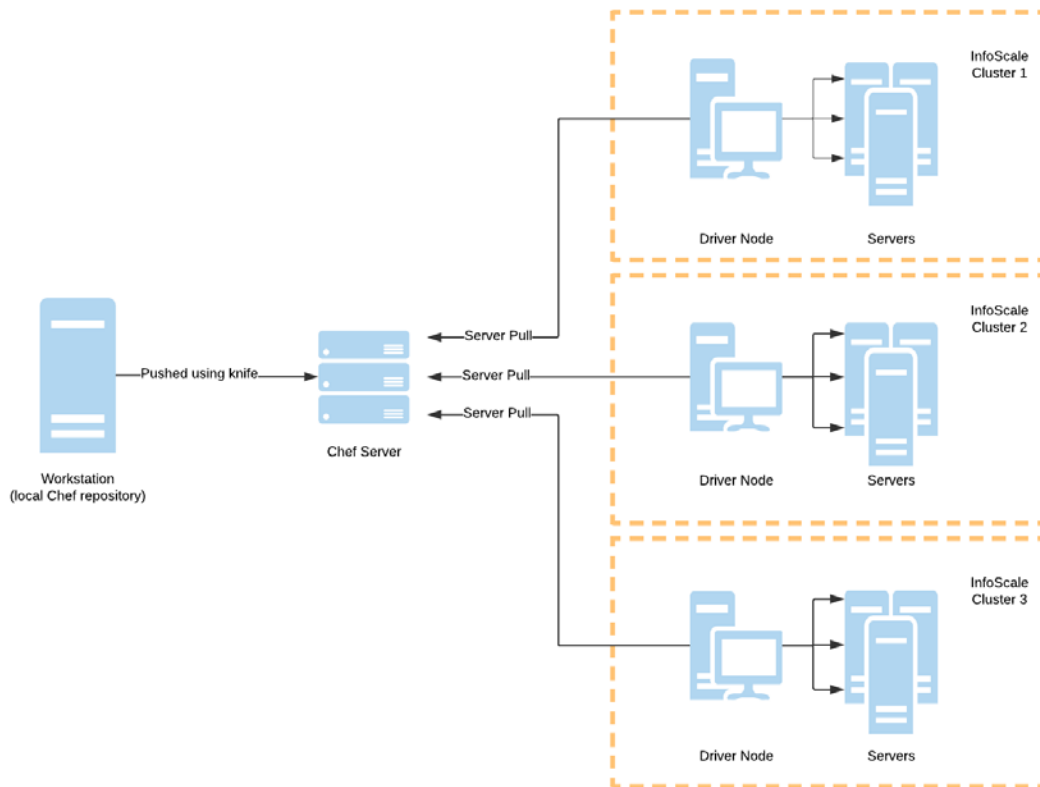
The following platforms are supported by the *veritas\_infoscale* cookbook:

- Linux
- Solaris
- Aix

## Chef environment

Ensure that you have installed the Chef infrastructure in your environment. A typical Chef environment usually consists of the following components:

- **Workstation:** The Chef development kit is set up on this server. The workstation server is used to develop and test cookbooks and recipes.
- **Chef server:** The Chef server stores cookbooks, node policies, and node metadata, which are accessed and used by the Chef client on the Chef nodes.
- **Chef node:** A Chef node is used to run the Chef client. Each cluster must contain a single Chef node, which acts as a driver node for the cluster.
- **Knife:** A command-line tool that provides an interface between a local chef repository (workstation) and the Chef server.



Refer to [Chef documentation](#) for instructions on installing and setting up the Workstation, Chef server, Chef node, and Knife components in your environment.

## Using Chef to install and configure InfoScale in a Linux, Solaris, or AIX environment

This guide provides information about using the `veritas_infoscale` cookbook. To perform any operation in Veritas InfoScale using Chef, you must perform the following high-level steps:

### [Prerequisites](#)

[Step 1: Downloading the Veritas InfoScale cookbook from the supermarket](#)

[Step 2: Specifying operations in a recipe file](#)

[Step 3: Creating roles and defining attributes](#)

[Step 4: Running the Chef-client on the Chef node](#)

Note that you will require basic knowledge of the *JavaScript Object Notation (JSON)* format and *Chef* while performing the procedure in this guide.

## Prerequisites

Ensure that the following prerequisites are met in your environment:

- Passwordless communication is established between all servers in a cluster.  
**Tip:** The user can use the *pl* utility to set up the ssh and rsh connections automatically.
- Veritas InfoScale installation files are downloaded on the chef node where chef-client is run.
- (Only for Veritas InfoScale version 7.4 and later) All required Veritas InfoScale license files (with a .slf extension) are accessible from the local server. Contact the Customer Care of your region to procure an applicable slf license key file. Refer to the following link for contact information of the Customer Care center for your region:  
[https://www.veritas.com/content/support/en\\_US/contact-us.html](https://www.veritas.com/content/support/en_US/contact-us.html)
- The veritas\_infoscale cookbook uses response files to perform all operations such as installation, upgrade, patch upgrade, and so on. Before using the veritas\_infoscale cookbook, ensure that you have read and understood all prerequisites that are applicable while performing those operations with response files. For more information about using response files, refer to the following guides from the [Veritas Documentation Library](#):
  - *Veritas InfoScale Installation Guide*
  - *Storage Foundation Configuration and Upgrade Guide*
  - *Cluster Server Configuration and Upgrade Guide*
  - *Storage Foundation and High Availability Configuration and Upgrade Guide*
  - *Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide*
  - *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*
  - *Storage Foundation for Sybase ASE CE Configuration and Upgrade Guide*
- Ensure that you read through and agree with the End User License Agreement in the PDF file (located in the installation media), before proceeding with any operation in Veritas InfoScale using Chef.

## Step 1: Downloading the Veritas InfoScale cookbook from the supermarket

You can download the veritas\_infoscale cookbook from the Chef supermarket to use its resources.

### To download the Veritas InfoScale cookbook from the public Supermarket

1. To see a list of all community cookbooks available from Supermarket, run the following:

```
knife cookbook site list
```

2. Search for the veritas\_infoscale cookbook using the following command:

```
knife cookbook site search veritas_infoscale
```

3. Download the veritas\_infoscale cookbook using the following command:

```
knife cookbook site download veritas_infoscale
```

The cookbook is downloaded in a tar.gz package. Use the following command to extract the package: `tar xvzf veritas_infoscale.tar.gz`

If you are in an air-gapped environment you can download the cookbook from the [Chef Supermarket](#) website and physically transfer the cookbook to your local network using removable storage. Alternatively, you can refer to the Chef documentation for steps on setting up a private supermarket in your local network.

**Important:** After you have downloaded and extracted the veritas\_infoscale cookbook, ensure that you upload the cookbook from the current working directory in the chef repository to the Chef server. Use the following command:

```
knife cookbook upload veritas_infoscale
```

For more information about uploading cookbooks, see the [Chef documentation](#).

## Step 2: Specifying operations in a recipe file

Before you run the Chef client on a Chef node, you must specify the operations you want to perform in the recipe file you are using. The following resources are shipped in the veritas\_infoscale cookbook:

- `infoscale_deploy`
- `infoscale_configure`

Depending on which operations you want to perform, enter the required values in the action block within the `infoscale_deploy` or `infoscale_configure` resource.

Value	Resource name	Description
Install	<code>infoscale_deploy</code>	Installs Veritas InfoScale products in your environment
Configure	<code>infoscale_deploy</code>	Configures the various components used by Veritas InfoScale products
fullupgrade	<code>infoscale_deploy</code>	Upgrades Veritas InfoScale products
patch_upgrade	<code>infoscale_deploy</code>	Upgrades Veritas InfoScale products to a particular patch
Uninstall	<code>infoscale_deploy</code>	Uninstalls Veritas InfoScale products
rollingupgrade	<code>infoscale_deploy</code>	Performs a rolling upgrade for highly available clusters
license	<code>infoscale_deploy</code>	Manages Veritas InfoScale licenses

security	infoscale_configure	Configures security in Veritas InfoScale environment
fencing	infoscale_configure	Configures I/O fencing to prevent data corruption in the event of a communication failure. Note that only disk-based, majority-based, and disabled-based fencing can be configured using Chef

## Syntax:

```
<resource_name> '<instance_name>' do
  action [ :<value_1> :<value_2> :<value_n>]
end
```

Replace the following variables in the above syntax:

- `<resource_name>` is the name of the resource used by the operation.
- `<value_1>`, `<value_2>`, ... `<value_n>` are n number of values corresponding to the operations you want to perform. The operations will be performed in the sequence in which the values are entered.
- `<instance_name>` is the name of the resource instance. You can provide any instance name for the resource.

Because we will only be installing and configuring InfoScale, we will use the **infoscale\_deploy** resource in our recipe, as follows:

```
# Cookbook:: veritas_infoscale
# Recipe:: default
#
# Copyright:: 2017, The Authors, All Rights Reserved.

infoscale_deploy 'Infoscale' do
  action [ :install, :configure ]
end
~
~
~
~
```

Ensure that you upload the cookbook from the current working directory in the chef repository to the Chef server. Use the following command:

```
knife cookbook upload veritas_infoscale
```

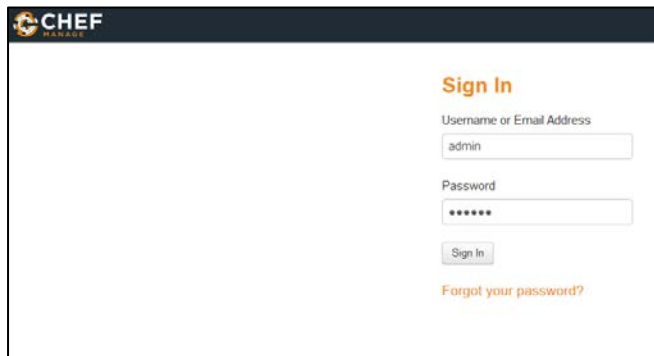
## Step 3: Creating roles and defining attributes

Chef roles can be used to perform operations to manage cluster configurations in a Veritas InfoScale environment. A separate role (specifying its run list and override attributes) must be created for each cluster

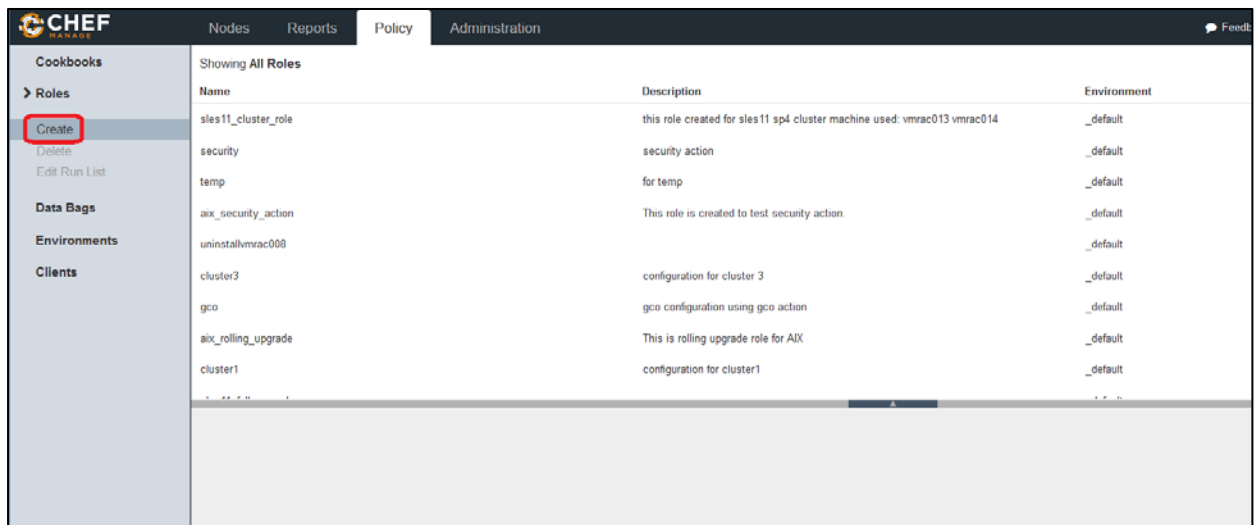
configuration in your environment. When a role is run against a node, the configuration details of that node are compared against the attributes of the role, and then the contents of that role's run-list are applied to the node's configuration details.

To create a role for installation and configuration using the Chef Manage interface

1. Log on to the Chef server with Administrator privileges.



2. Click on the Policy tab at the top of the page.
3. Expand Role in the left pane and click Create. The Create Role dialog box appears.



4. Enter a name and description to help you easily identify the role.



**Create Role** [X]

Basics ▶ Run List ▶ Default Attributes ▶ Override Attributes

Name  
install\_and\_configure

Description  
Use this Role to install and configure InfoScale in your environment.

Cancel < Previous Next > Create Role

5. Click Next.
6. Search for the recipe that you modified in [Step 2: Specifying operations in a recipe file](#) in the Available Recipes list.
7. Drag and drop the recipe to the Create Run List box so that the recipe is run whenever the Chef client is invoked using this role. Ensure that you have customized the default recipe file to perform the desired operations in your environment (see [Step 2: Specifying operations in a recipe file](#)).

**Create Role** [X]

Basics ▶ Run List ▶ Default Attributes ▶ Override Attributes

Available Roles

- sles11\_cluster\_role
- security
- temp
- aix\_security\_action
- uninstallvmrac008
- cluster3

Available Recipes

- depends\_test
- veritas\_infoscale\_test
- veritas\_infoscale\_test::default1
- veritas\_infoscale\_test::default2
- veritas\_infoscale\_test::default\_aix
- veritas\_infoscale\_test::default\_fencing

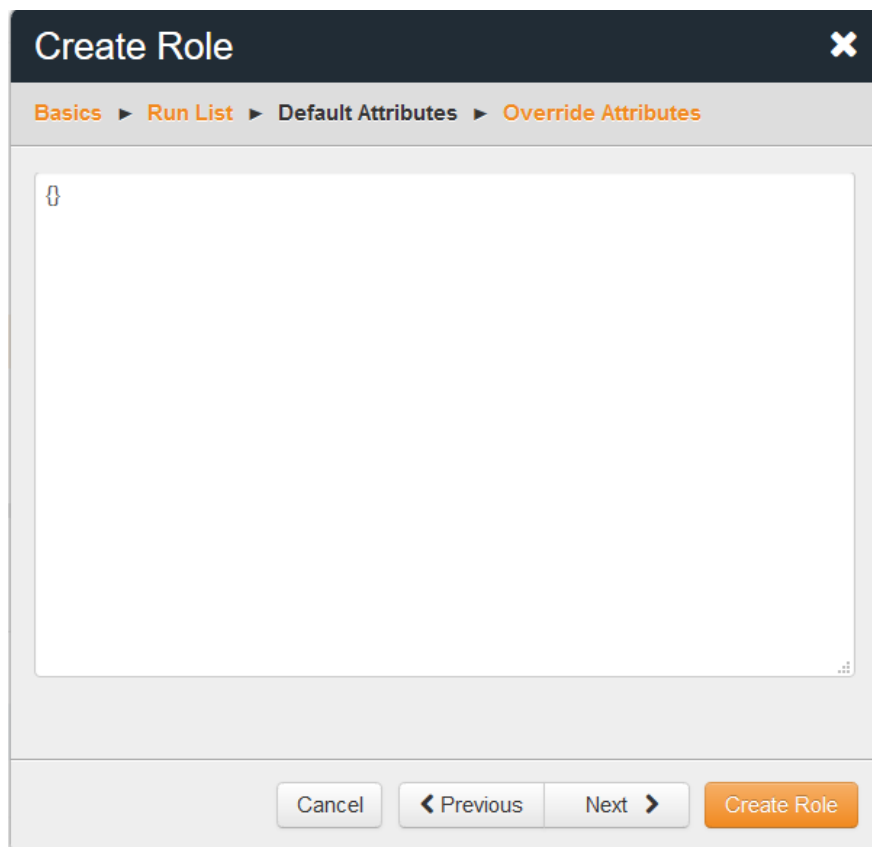
Current Run List

- veritas\_infoscale\_test::install\_configure

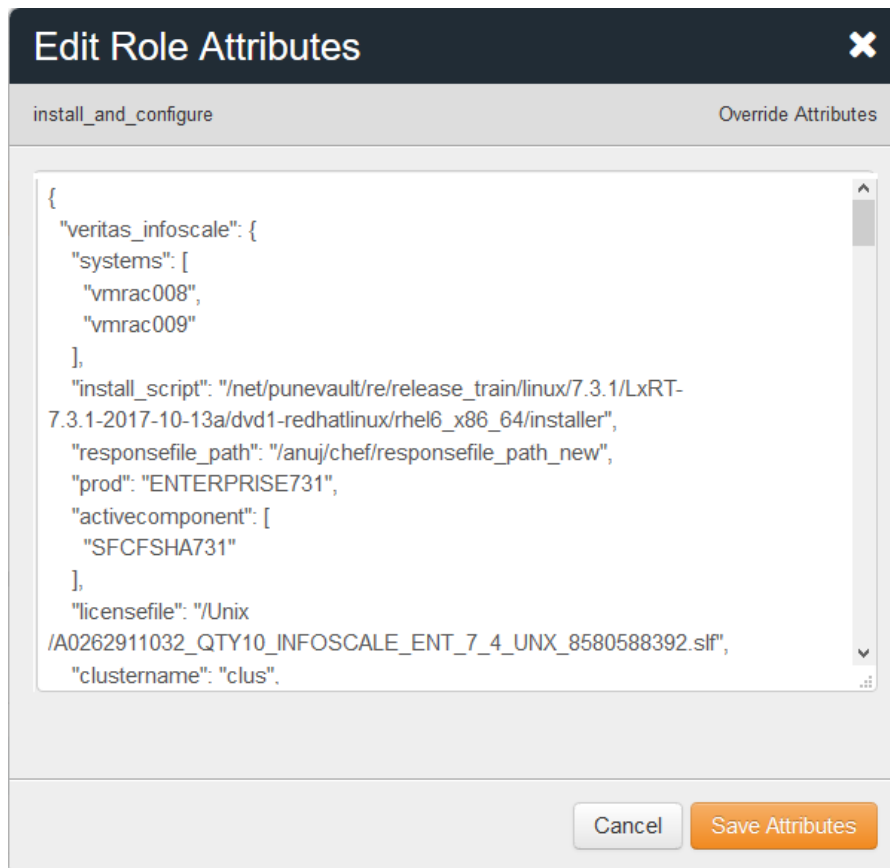
Cancel < Previous Next > Create Role

8. Click Next.
9. You do not need to enter default attributes in the role, because the default values for Veritas InfoScale

operations are pre-defined in the attribute file of the default.rb recipe.



10. Click Next.
11. Enter the override attributes for the role in JSON format. Refer to the [Attributes for installing and configuring InfoScale](#) section for a list of the attributes used in a Veritas InfoScale operation. You can also refer to [Examples of attributes in JSON format](#) for examples of attributes based on the operation you want to perform.



12. Click Create Role.

Defining attributes in a role

The attributes can be of the following two types:

**Optional:** Optional attributes need not be defined by the user. The default values of optional attributes are pre-defined with default values in the attribute file of the default.rb recipe. However, you can override the default values by defining override type of attributes at the role-level.

**Mandatory:** All mandatory attributes required in an operation must be defined at the role-level as an override type of attribute.

**Important:** The veritas\_infoscale cookbook uses response files to perform all operations such as installation, upgrade, patch upgrade, and so on. Before using the veritas\_infoscale cookbook, ensure that you have read and understood all prerequisites that are applicable while performing those operations with response files. For more information about using response files, refer to the following guides:

- *Installation Guide*
- *Storage Foundation Configuration and Upgrade Guide*
- *Cluster Server Configuration and Upgrade Guide*
- *Storage Foundation and High Availability Configuration and Upgrade Guide*
- *Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide*

- *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*
- *Storage Foundation for Sybase ASE CE Configuration and Upgrade Guide*

Each cluster in the Chef-InfoScale environment consists of a Chef node. Veritas InfoScale replicates the configuration of the Chef node to other nodes in the cluster. In this way, we can use role attributes to define the configuration of the Chef node, and subsequently define the configuration of the cluster.

Attributes for installing and configuring InfoScale

The attributes listed in the below table are used while performing an installation and configuration of Veritas InfoScale products. Ensure that you specify values for all mandatory attributes in the role you are creating. Optional attributes which are not specified by the user will use their pre-defined default values.

Attribute	Dimension	Type	Description
systems	List	Mandatory	<p>List of the systems on which the product is to be installed and configured.</p> <p><b>Syntax</b></p> <pre>"systems": ["&lt;server_1&gt;", "&lt;server_2&gt;", ... , "&lt;server_n&gt;"]</pre> <p>&lt;server_1&gt;, &lt;server_2&gt;, ...&lt;server_n&gt; are the host names or IP addresses of n number of systems on which you want to perform the installation.</p> <p><b>Example</b></p> <pre>"systems": ["vmrac008", "vmrac009"]</pre>
install_script	Scalar	Mandatory	<p>Path to where the Veritas InfoScale installation files are stored.</p> <p>For example, <b>&lt;temporary_location&gt;/installer.</b></p>
prod	Scalar	Mandatory	<p>Specifies which Veritas InfoScale product and version you want to install.</p> <p><b>Syntax</b></p> <pre>&lt;product_name&gt;&lt;version_number&gt;</pre> <p>&lt;product_name&gt; can be either ENTERPRISE, AVAILABILITY, STORAGE, or FOUNDATION.</p> <p>&lt;version_number&gt; is the version number of the Veritas InfoScale product.</p> <p><b>Examples</b></p> <pre>ENTERPRISE74, AVAILABILITY731, STORAGE70, or FOUNDATION70</pre>
keyless	List	Mandatory ( <i>Even though keyless is a</i>	<p>List of the product keys for keyless installation that are to be registered on the systems</p>

Attribute	Dimension	Type	Description
		<i>mandatory attribute, you must leave the keyless attribute empty if you are providing values for either the license or licensefile attribute.)</i>	specified in the systems attribute. The value can be either ENTERPRISE, AVAILABILITY, STORAGE, or FOUNDATION.
license	List	Mandatory ( <i>Even though license is a mandatory attribute, you must leave the license attribute empty if you are providing values for either the keyless or licensefile attribute.</i> )	List of license keys to be registered on the systems specified in the systems attribute.
licensefile (only for versions 7.4 and later)	List	Mandatory ( <i>Even though licensefile is a mandatory attribute, you must leave the licensefile attribute empty if you are providing values for either the keyless or license attribute.</i> )	List of paths to the slf license file to be registered on the systems specified in the systems attribute. Ensure that the license files are accessible from the server where the installer is run.
clustername	Scalar	Mandatory	Defines the name of the cluster.

Attribute	Dimension	Type	Description
clusterid	Scalar	Optional ( <i>If value not provided, we randomly generate number and assign it.</i> )	Defines an integer value between 0 and 65535 that uniquely identifies the cluster.
heartbeat_links_interface	List	Mandatory	<p>Defines the NIC to be used for a private heartbeat link on each system. At least two LLT links are required per system (lltlink1 and lltlink2).</p> <p><b>Syntax 1:</b></p> <pre>"heartbeat_links_interface": {   "&lt;server_1&gt;": [     "&lt;NIC_1&gt;",     "&lt;NIC_2&gt;",     "&lt;NIC_3&gt;",     "&lt;NIC_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;NIC_1&gt;",     "&lt;NIC_2&gt;",     "&lt;NIC_3&gt;",     "&lt;NIC_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;NIC_1&gt;",     "&lt;NIC_2&gt;",     "&lt;NIC_3&gt;",     "&lt;NIC_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;NIC_1&gt;",     "&lt;NIC_2&gt;",     "&lt;NIC_3&gt;",     "&lt;NIC_4&gt;"   ] }</pre>

Attribute	Dimension	Type	Description
			<p>]</p> <ul style="list-style-type: none"> <li>• <code>&lt;server_1&gt;</code>, <code>&lt;server_2&gt;</code>, ... <code>&lt;server_n&gt;</code> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <code>&lt;NIC_1&gt;</code>, <code>&lt;NIC_2&gt;</code>, ... <code>&lt;NIC_4&gt;</code> are the NICs that each private heartbeat link uses on that particular system. You must enter at least 2 entries per system, one for each heartbeat link.</li> </ul> <p><b>Syntax 2:</b></p> <pre>"heartbeat_links_interface": ["&lt;NIC_1&gt;","&lt;NIC_2&gt;","&lt;NIC_3&gt;","&lt;NIC_4&gt;"]</pre> <ul style="list-style-type: none"> <li>• <code>&lt;NIC_1&gt;</code>, <code>&lt;NIC_2&gt;</code>, ... <code>&lt;NIC_4&gt;</code> are the NICs that each private heartbeat link uses on all systems in the cluster. You must enter at least 2 entries per system, one for each heartbeat link.</li> </ul>
lopri_link_interface	List	Optional	<p>Lists the NIC interfaces to be configured for low priority heartbeat links.</p> <p>Typically, <code>lopri_link_interface</code> is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, <code>lltlinklowpri1</code>, <code>lltlinklowpri2</code>, and so on. You must enclose the system name within double quotes. It is recommended to configure at least one low priority link in the cluster.</p> <p><b>Syntax 1:</b></p> <pre>"lopri_link_interface": {</pre>



Attribute	Dimension	Type	Description
			<pre> "&lt;server_1&gt;": [   "&lt;NIC_1&gt;",   "&lt;NIC_2&gt;",   "&lt;NIC_3&gt;",   "&lt;NIC_4&gt;" ], "&lt;server_2&gt;": [   "&lt;NIC_1&gt;",   "&lt;NIC_2&gt;",   "&lt;NIC_3&gt;",   "&lt;NIC_4&gt;" ], "&lt;server_3&gt;": [   "&lt;NIC_1&gt;",   "&lt;NIC_2&gt;",   "&lt;NIC_3&gt;",   "&lt;NIC_4&gt;" ], "&lt;server_n&gt;": [   "&lt;NIC_1&gt;",   "&lt;NIC_2&gt;",   "&lt;NIC_3&gt;",   "&lt;NIC_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;NIC_1&gt;, &lt;NIC_2&gt;, ... &lt;NIC_4&gt; are the NICs that each private low priority heartbeat link uses on that particular system.</li> </ul> <p><b>Syntax 2:</b></p> <pre> "lopri_link_interface": ["&lt;NIC_1&gt;","&lt;NIC_2&gt;","&lt;NIC_3&gt;","&lt;NIC_4&gt; </pre>

Attribute	Dimension	Type	Description
			<p>"]</p> <ul style="list-style-type: none"> <li>&lt;NIC_1&gt;, &lt;NIC_2&gt;, ... &lt;NIC_4&gt; are the NICs that each private low priority heartbeat link uses on all systems in the cluster.</li> </ul>
lltoverudp	Scalar	Optional	Specifies whether to configure heartbeat link using LLT over UDP. Default value is 0.
heartbeat_udplink_port	Hash (two-dimensional array)	Mandatory ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	<p>Lists the UDP port number (16-bit integer value) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the UDP port number for each of the heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre>"heartbeat_udplink_port": {   "&lt;server_1&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ], }</pre>

Attribute	Dimension	Type	Description
			<pre>"&lt;server_n&gt;": [   "&lt;port_1&gt;",   "&lt;port_2&gt;",   "&lt;port_3&gt;",   "&lt;port_4&gt;" ]</pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;port_1&gt;, &lt;port_2&gt;, ... &lt;port_4&gt; are the port numbers that each heartbeat link uses on that particular system.</li> </ul>
heartbeat_udplinklowpri_port	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	<p>Lists the UDP port number (16-bit integer value) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the UDP port number for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre>"heartbeat_udplinklowpri_port": {   "&lt;server_1&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ] }</pre>

Attribute	Dimension	Type	Description
			<pre> ], "&lt;server_3&gt;": [   "&lt;port_1&gt;",   "&lt;port_2&gt;",   "&lt;port_3&gt;",   "&lt;port_4&gt;" ], "&lt;server_n&gt;": [   "&lt;port_1&gt;",   "&lt;port_2&gt;",   "&lt;port_3&gt;",   "&lt;port_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• <i>&lt;server_1&gt;</i>, <i>&lt;server_2&gt;</i>, ... <i>&lt;server_n&gt;</i> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <i>&lt;port_1&gt;</i>, <i>&lt;port_2&gt;</i>, ... <i>&lt;port_4&gt;</i> are the port numbers that each low priority heartbeat link uses on that particular system.</li> </ul>
heartbeat_udplink_address	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	<p>Lists the IP address (IPv4 or IPv6) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the IP addresses for each of the heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_udplink_address": {   "&lt;server_1&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;" </pre>

Attribute	Dimension	Type	Description
			<pre> ], "&lt;server_2&gt;": [   "&lt;IP_address_1&gt;",   "&lt;IP_address_2&gt;",   "&lt;IP_address_3&gt;",   "&lt;IP_address_4&gt;" ], "&lt;server_3&gt;": [   "&lt;IP_address_1&gt;",   "&lt;IP_address_2&gt;",   "&lt;IP_address_3&gt;",   "&lt;IP_address_4&gt;" ], "&lt;server_n&gt;": [   "&lt;IP_address_1&gt;",   "&lt;IP_address_2&gt;",   "&lt;IP_address_3&gt;",   "&lt;IP_address_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;IP_address_1&gt;, &lt;IP_address_2&gt;, ... &lt;IP_address_4&gt; are the IP addresses that each heartbeat link uses on that particular system.</li> </ul>
heartbeat_udplink_netmask	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	Lists the netmask (prefix for IPv6) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the netmask for each of the heartbeat links are specified in a hash (two dimensional array) in JSON.

Attribute	Dimension	Type	Description
			<p><b>Syntax:</b></p> <pre> "heartbeat_udplink_netmask": {   "&lt;server_1&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ] } </pre> <ul style="list-style-type: none"> <li>• <i>&lt;server_1&gt;</i>, <i>&lt;server_2&gt;</i>, ... <i>&lt;server_n&gt;</i> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <i>&lt;netmask_1&gt;</i>, <i>&lt;netmask_2&gt;</i>, ... <i>&lt;netmask_4&gt;</i> are the netmasks that each heartbeat link uses on that particular system.</li> </ul>

Attribute	Dimension	Type	Description
heartbeat_udplinklowpri_address	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	<p>Lists the IP address (IPv4 or IPv6) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the IP addresses for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_udplinklowpri_address": {   "&lt;server_1&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ] } </pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP</li> </ul>

Attribute	Dimension	Type	Description
			<p>addresses of n number of systems in the cluster.</p> <ul style="list-style-type: none"> <li>• <i>&lt;IP_address_1&gt;</i>, <i>&lt;IP_address_2&gt;</i>, ... <i>&lt;IP_address_4&gt;</i> are the IP addresses that each low priority heartbeat link uses on that particular system.</li> </ul>
heartbeat_udplinklowpri_netmask	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over UDP.</i> )	<p>Lists the netmask (prefix for IPv6) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the netmask for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_udplinklowpri_netmask": {   "&lt;server_1&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;netmask_1&gt;", </pre>



Attribute	Dimension	Type	Description
			<pre>"&lt;netmask_2&gt;", "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ]</pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;netmask_1&gt;, &lt;netmask_2&gt;, ... &lt;netmask_4&gt; are the netmasks that each low priority heartbeat link uses on that particular system.</li> </ul>
lltoverrdma	Scalar	Mandatory	Indicates whether to configure heartbeat link using LLT over Remote Direct Memory Access (RDMA).
heartbeat_rdmalink_address	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over RDMA.</i> )	<p>Lists the IP address (IPv4 or IPv6) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the IP addresses for each of the heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre>"heartbeat_rdmalink_address": {   "&lt;server_1&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",</pre>

Attribute	Dimension	Type	Description
			<pre> "&lt;IP_address_4&gt;" ], "&lt;server_3&gt;": [   "&lt;IP_address_1&gt;",   "&lt;IP_address_2&gt;",   "&lt;IP_address_3&gt;",   "&lt;IP_address_4&gt;" ], "&lt;server_n&gt;": [   "&lt;IP_address_1&gt;",   "&lt;IP_address_2&gt;",   "&lt;IP_address_3&gt;",   "&lt;IP_address_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;IP_address_1&gt;, &lt;IP_address_2&gt;, ... &lt;IP_address_4&gt; are the IP addresses that each heartbeat link uses on that particular system.</li> </ul>
heartbeat_rdmalink_netmask	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over RDMA.</i> )	<p>Lists the netmask (prefix for IPv6) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the netmask for each of the heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_rdmalink_netmask": {   "&lt;server_1&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;", </pre>

Attribute	Dimension	Type	Description
			<pre> "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ], "&lt;server_2&gt;": [ "&lt;netmask_1&gt;", "&lt;netmask_2&gt;", "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ], "&lt;server_3&gt;": [ "&lt;netmask_1&gt;", "&lt;netmask_2&gt;", "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ], "&lt;server_n&gt;": [ "&lt;netmask_1&gt;", "&lt;netmask_2&gt;", "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• <i>&lt;server_1&gt;</i>, <i>&lt;server_2&gt;</i>, ... <i>&lt;server_n&gt;</i> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <i>&lt;netmask_1&gt;</i>, <i>&lt;netmask_2&gt;</i>, ... <i>&lt;netmask_4&gt;</i> are the netmasks that each heartbeat link uses on that particular system.</li> </ul>
heartbeat_rdmalink_port	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over RDMA.</i> )	Lists the RDMA port number (16-bit integer value) that each heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the RDMA port number (16-bit integer value) for each of the heartbeat

Attribute	Dimension	Type	Description
			<p>links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_rdmalink_port": {   "&lt;server_1&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ] } </pre> <ul style="list-style-type: none"> <li>• <i>&lt;server_1&gt;</i>, <i>&lt;server_2&gt;</i>, ... <i>&lt;server_n&gt;</i> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <i>&lt;port_1&gt;</i>, <i>&lt;port_2&gt;</i>, ... <i>&lt;port_4&gt;</i> are the port numbers that each</li> </ul>

Attribute	Dimension	Type	Description
			heartbeat link uses on that particular system.
heartbeat_rdmalinklowpri_address	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over RDMA.</i> )	<p>Lists the IP address (IPv4 or IPv6) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the IP addresses for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_rdmalinklowpri_address": {   "&lt;server_1&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;IP_address_1&gt;",     "&lt;IP_address_2&gt;",     "&lt;IP_address_3&gt;",     "&lt;IP_address_4&gt;"   ] } </pre>

Attribute	Dimension	Type	Description
			<p>&lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster. &lt;IP_address_1&gt;, &lt;IP_address_2&gt;, ... &lt;IP_address_4&gt; are the IP addresses that each low priority heartbeat link uses on that particular system.</p>
heartbeat_rdmalinklowpri_netmask	Hash (two-dimensional array)	Optional ( <i>This attribute is mandatory if LLT is configured over RDMA.</i> )	<p>Lists the netmask (prefix for IPv6) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the netmask for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre> "heartbeat_rdmalinklowpri_netmask": {   "&lt;server_1&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_3&gt;": [     "&lt;netmask_1&gt;",     "&lt;netmask_2&gt;",     "&lt;netmask_3&gt;",     "&lt;netmask_4&gt;"   ],   "&lt;server_n&gt;": [     "&lt;netmask_1&gt;", </pre>

Attribute	Dimension	Type	Description
			<pre>"&lt;netmask_2&gt;", "&lt;netmask_3&gt;", "&lt;netmask_4&gt;" ]</pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;netmask_1&gt;, &lt;netmask_2&gt;, ... &lt;netmask_4&gt; are the netmasks that each low priority heartbeat link uses on that particular system.</li> </ul>
<p>heartbeat_rdmalinklowpri_port</p>	<p>Hash (two-dimensional array)</p>	<p>Optional (<i>This attribute is mandatory if LLT is configured over RDMA.</i>)</p>	<p>Lists the RDMA port number (16-bit integer value) that each low priority heartbeat link uses on each of the nodes in the cluster. The system name of each of the nodes and the RDMA port number for each of the low priority heartbeat links are specified in a hash (two dimensional array) in JSON.</p> <p><b>Syntax:</b></p> <pre>"heartbeat_rdmalinklowpri_port": {   "&lt;server_1&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_2&gt;": [     "&lt;port_1&gt;",     "&lt;port_2&gt;",     "&lt;port_3&gt;",     "&lt;port_4&gt;"   ],   "&lt;server_3&gt;": [</pre>

Attribute	Dimension	Type	Description
			<pre> "&lt;port_1&gt;", "&lt;port_2&gt;", "&lt;port_3&gt;", "&lt;port_4&gt;" ], "&lt;server_n&gt;": [   "&lt;port_1&gt;",   "&lt;port_2&gt;",   "&lt;port_3&gt;",   "&lt;port_4&gt;" ] </pre> <ul style="list-style-type: none"> <li>• <i>&lt;server_1&gt;</i>, <i>&lt;server_2&gt;</i>, ... <i>&lt;server_n&gt;</i> are the host names or IP addresses of n number of systems in the cluster.</li> <li>• <i>&lt;port_1&gt;</i>, <i>&lt;port_2&gt;</i>, ... <i>&lt;port_4&gt;</i> are the port numbers that each low priority heartbeat link uses on that particular system.</li> </ul>
userenpw	List	Optional	<p>List of encoded passwords for users. The value in the list can be "Administrators Operators Guests"</p> <p><b>Note:</b> The order of the values for the userenpw list must match the order of the values in the username list.</p>
username	List	Optional	Lists the names of users.
userpriv	List	Optional	<p>Lists the privileges to be granted to the users.</p> <p><b>Note:</b> The order of the values for the userpriv list must match the order of the values in the username list.</p>
eat_security	Scalar	Optional	Specifies whether the cluster should be enabled with the secure mode or not. By default, this



Attribute	Dimension	Type	Description
			value is set to 1, which means that security will be enabled in the cluster.
eat_security_fips	Scalar	Optional	Specifies whether to enable or disable security with FIPS mode on a running VCS cluster. By default, this value is set to 0, which means that security with FIPS mode will be disabled in the cluster.
keyfile	Scalar	Optional	Location of the SSH key file that is used by the Chef node to communicate with other nodes in the cluster.
logpath	Scalar	Optional	Specifies the location where the log files are to be stored. If no value is defined the default location is <b>/opt/VRTS/install/logs</b> .
noipc	Scalar	Optional	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates. (0: sends outbound networking calls to SORT, 1: does not send outbound networking calls to SORT). The default value is 0.
patch_path	List ( <i>maximum 5 patches</i> )	Optional	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
responsefile_path	Scalar	Optional	Temporary location where you want to store the response file created during installation. By default, the response file is stored in the home directory of the user. You can override the default path by specifying a different path for this

Attribute	Dimension	Type	Description
			attribute. Note that after successful installation is complete the response file is moved to <b>/opt/VRTS/install/chef</b> .
rsh	Scalar	Optional	Specifies that RSH protocol must be used instead of SSH protocol while communicating between systems. A boolean value of 1 indicates that RSH will be used and a boolean value of 0 indicates that SSH will be used. This is an optional attribute and if no value is provided SSH protocol is used by default.
tmppath	Scalar	Optional	Specifies the location where temporary files and dependent RPMs are stored during the install. The default location is <b>/var/tmp</b> .
uploadlogs	Scalar	Optional	Specifies whether you want to upload logs to the Veritas website. The value 1 indicates that the installation logs are uploaded to the Veritas website. The value 0 indicates that the installation logs are not uploaded to the Veritas website. By default this value is set to 1.
activecomponent	List	Mandatory	Specifies the component for operations like precheck, configure, addnode, install, and configure (together). The value can be either of the following: <ul style="list-style-type: none"> <li>- VCS&lt;VersionNumber&gt; for VCS component.</li> <li>- SF&lt;VersionNumber&gt; for SF component.</li> <li>- SVS&lt;VersionNumber&gt; for SVS component.</li> <li>- SFRAC&lt;VersionNumber&gt; for SF Oracle RAC component.</li> <li>- SFSYBASECE&lt;VersionNumber&gt; for</li> </ul>

			<p>SFSYBASECE component.</p> <ul style="list-style-type: none"> <li>- SFCFSHA&lt;VersionNumber&gt; for SFCFSHA component.</li> <li>- SFHA&lt;VersionNumber&gt; for SFHA component.</li> </ul>
fencing_option	Scalar	Optional	<p>Specifies the I/O fencing configuration mode. The default value is 1.</p> <ul style="list-style-type: none"> <li>• 1: Disk-based fencing</li> <li>• 2: Majority-based fencing</li> <li>• 3: Disabled-based I/O fencing</li> </ul>
fencingenabled	Scalar	Optional	<p>Specifies whether to enable fencing in your Veritas product. Valid values are 0 or 1. The default value is 0.</p>
fencing_auto_refresh_reg	Scalar	Optional	<p>Enable the auto refresh of coordination points variable in case registration keys are missing on any of CP servers. The default value is 0.</p>
fencing_config_cpagent	Scalar	Optional	<p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent. The default value is 1.</p>
fencing_cpagent_monitor_freq	Scalar	Optional	<p>Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution. <b>Note:</b> Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidentally deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk</p>

			Group constitution every five monitor cycles. If LevelTwoMonitorFreq attribute is not set, the agent will not monitor any changes to the Coordinator Disk Group. 0 means not to monitor the Coordinator Disk Group constitution. The default value is 5
fencing_cpagentgrp	Scalar	Optional	Name of the service group which will have the Coordination Point agent resource as part of it. The default value is "vxfen".  <b>Note:</b> This field is obsolete if the fencing_config_cpagent field is given a value of '0'.
fencing_dgname	Scalar	Optional <i>(Mandatory if fencingenabled is 1 and fencing_option is set to 1)</i>	Specifies the disk group for I/O fencing. <b>Note:</b> You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
fencing_newdg_disks	List	Optional <i>(Mandatory if fencingenabled is 1 and fencing_option is set to 1)</i>	Specifies the disks to use to create a new disk group for I/O fencing. <b>Note:</b> You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
allowcomms	Scalar	Optional	Indicates whether to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). By default, this value is set to 1.
defaultaccess	Scalar	Optional	Specifies whether to grant read access to all. By default this value is set to 0.
donotreconfigurefencing	Scalar	Optional	Specifies whether you want to re-configure fencing. Note that any customizations made to the configuration will be deleted and the default

			configuration will be re-installed. By default this value is set to 1.
donotreconfigurevcs	Scalar	Optional	Specifies whether you want to re-configure VCS. Note that any customizations made to the configuration will be deleted and the default configuration will be re-installed. By default this value is set to 1.
secusrgrps	List	Optional	Specifies the user groups which will receive read access to the cluster.
vm_no_open_vols	Scalar	Optional	Specifies whether to ask the user if there are any open volumes (when vxconfigd is not enabled). Such prompts are asked during uninstallations. (1: affirms there are no open volumes on the system). By default value is set to 0.
csgnetmask	Scalar	Optional ( <i>This attribute must be specified if (Cluster Service Group) CSG is used.</i> )	Defines the netmask of the virtual IP address for the cluster.
csgnic	Scalar	Optional ( <i>This attribute must be specified if (Cluster Service Group) CSG is used.</i> )	Defines the NIC device to use on a system. <b>Syntax 1:</b> <pre>"csgnic": {   "&lt;server_1&gt;": [     "&lt;NIC&gt;"   ],   "&lt;server_2&gt;": [     "&lt;NIC&gt;"   ],   "&lt;server_3&gt;": [     "&lt;NIC&gt;"   ],   "&lt;server_n&gt;": [</pre>

			<pre>"&lt;NIC&gt;" ]</pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;NIC&gt; is the NIC that is used on that particular system.</li> </ul> <p><b>Syntax 2:</b></p> <pre>"csgnic": "&lt;NIC&gt;"</pre> <ul style="list-style-type: none"> <li>• &lt;NIC&gt; is the NIC that is used on all the systems.</li> </ul>
csgvip	Scalar	Optional ( <i>This attribute must be specified if (Cluster Service Group) CSG is used.</i> )	Defines the virtual IP address for the cluster.
gconetmask	Scalar	Optional ( <i>This attribute must be specified if GCO is used.</i> )	Defines the Netmask of the virtual IP address that the Global Cluster Option uses.
gconic	Scalar	Optional ( <i>This attribute must be specified if GCO is used.</i> )	<p>Specifies the NIC for the virtual IP that the Global Cluster Option uses. You can enter all as a system value if the same NIC is used on all systems.</p> <p><b>Syntax 1:</b></p> <pre>"gconic": {   "&lt;server_1&gt;": [     "&lt;NIC&gt;"   ],   "&lt;server_2&gt;": [     "&lt;NIC&gt;"   ],   "&lt;server_3&gt;": [</pre>

			<pre>"&lt;NIC&gt;" ], "&lt;server_n&gt;": [   "&lt;NIC&gt;" ] </pre> <ul style="list-style-type: none"> <li>• &lt;server_1&gt;, &lt;server_2&gt;, ... &lt;server_n&gt; are the host names or IP addresses of n number of systems in the cluster.</li> <li>• &lt;NIC&gt; is the NIC that is used on that particular system.</li> </ul> <p><b>Syntax 2:</b></p> <pre>"gconic": "&lt;NIC&gt;"</pre> <ul style="list-style-type: none"> <li>• &lt;NIC&gt; is the NIC that is used on all the systems..</li> </ul>
gcovip	Scalar	Optional ( <i>This attribute must be specified if GCO is used.</i> )	Defines the virtual IP address that the Global Cluster Option uses.
smtprec	List	Optional ( <i>This attribute must be specified if SMTP protocol is used.</i> )	List of full email addresses (example: user@example.com) of SMTP recipients.
smtprsev	List	Optional ( <i>This attribute must be specified if SMTP protocol is used.</i> )	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients.
smtpserver	Scalar	Optional ( <i>This attribute must be specified if SMTP protocol is used.</i> )	Defines the domain-based host name (for example smtp.example.com) of the SMTP server to be used for web notification.
snmpcons	List	Optional ( <i>This attribute must be</i>	List of SNMP console system names.

		<i>specified if SNMP protocol is used.)</i>	
snmpcsev	List	Optional ( <i>This attribute must be specified if SNMP protocol is used.</i> )	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names.
snmpport	Scalar	Optional ( <i>This attribute must be specified if SNMP protocol is used.</i> )	Defines the SNMP trap daemon port (default is 162).

#### Step 4: Running the Chef-client on the Chef node

A Chef node is a node, which is managed by Chef, and is used to run the Chef-client. We recommend that you run the Chef client from a Chef node that is part of the cluster on which you are performing the operation. Use the following command to run the Chef-client on the Chef node:

```
chef-client -r "role[<name_of_the_role>]"
```

<name\_of\_the\_role> is the name of the role that you have created in [Step 3: Creating roles and defining attributes](#).

#### Example:

```
chef-client -r "role[install_and_configure]"
```

**Note:** Ensure that you have added a dependency on the veritas\_infoscale cookbook in your cookbook's metadata.rb, as follows:

```
depends 'veritas_infoscale'
```

#### Example of attributes used for installation and configuration

The following is an example of attributes entered (in JSON format) in a role for installing and configuring InfoScale:

```
{
  "veritas_infoscale": {
    "systems": [
```



```
    "vm1",
    "vm2"
  ],
  "install_script": "/mnt/7.2/dvd1-redhatlinux/rhel7_x86_64/installer ",
  "prod": "ENTERPRISE72",
  "activecomponent": [
    "SFHA72"
  ],
  "keyless": "ENTERPRISE",
  "license": "xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-xxxx-x",
  "clustertest": "clustertest",
  "heartbeat_links_interface": {
    "vm1": [
      "eth1",
      "eth2"
    ],
    "vm2": [
      "eth1",
      "eth2"
    ]
  },
  "lopri_link_interface": {
    "vm1": [
      "eth0"
    ],
    "vm2": [
      "eth0"
    ]
  },
  "username": [
    "user1",
    "user2"
  ],
  "userenpw": [
    "xxxxxxxxxxxx",
    "xxxxxxxxxxxx"
  ],
  "userpriv": "Administrators",
```

```
"smtpserver": "xxxxxxxxx.example.com",
"smtprecp": [
  "John.Johnson@example.com"
],
"smtpsev": [
  "Information"
],
"fencingenabled": "1",
"fencing_dgname": "vxfendg",
"fencing_option": "1",
}
}
```

## About Veritas Technologies LLC

Veritas Technologies LLC enables organizations to harness the power of their information, with solutions designed to serve the world's largest and most complex heterogeneous environments. Veritas works with 86 percent of Fortune 500 companies today, improving data availability and revealing insights to drive competitive advantage.

For specific country offices and contact numbers, please visit our website.

Veritas World Headquarters  
500 East Middlefield Road  
Mountain View, CA 94043  
+1 (650) 933 1000  
[www.veritas.com](http://www.veritas.com)

© 2018 Veritas Technologies LLC. All rights reserved. Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.