

# Mediant Cloud Edition (CE) SBC

## Deployment in Amazon AWS

Version 7.6



---

## Table of Contents

---

<b>1</b>	<b>Introduction .....</b>	<b>7</b>
1.1	Architecture Overview .....	7
1.2	Deployment Topology.....	8
1.3	Stack Manager .....	9
1.4	Deployment In a Single Availability Zone.....	9
1.5	Deployment In Multiple Availability Zone .....	11
1.5.1	Virtual IP Addresses .....	12
1.5.1.1	Using Virtual IP Addresses for Communication Across VPCs and with On-Prem Networks .....	14
1.5.2	AWS Network Load Balancer .....	17
1.5.2.1	Prerequisites for AWS Network Load Balancer Deployment .....	19
1.5.2.2	Enabling AWS Network Load Balancer.....	19
1.5.2.3	Configuration for AWS Network Load Balancer Deployment.....	20
<b>2</b>	<b>Installation Prerequisites.....</b>	<b>21</b>
2.1	Subscribing to AudioCodes Mediant VE Product in AWS Marketplace .....	21
2.2	IAM Role for Mediant CE .....	22
2.2.1	IAM Role for Initial Configuration from S3 URL .....	23
2.3	Cluster Subnet.....	24
2.3.1	Creating Cluster Subnet .....	24
2.3.2	Creating Private EC2 Endpoint in Cluster Subnet.....	26
2.3.3	Creating NAT Gateway in Cluster Subnet .....	28
2.4	Instance Types .....	30
<b>3</b>	<b>Deploying Mediant CE .....</b>	<b>31</b>
3.1	Public IP Addresses .....	37
3.2	Private IP Addresses for Deployments in Single Availability Zone .....	38
3.3	Private IP Addresses for Deployments in Multiple Availability Zones .....	40
3.4	Security Groups.....	41
3.4.1	Default Security Groups.....	41
3.4.2	Adjusting Default Security Groups.....	42
3.4.3	Using Custom Security Groups .....	43
3.5	Management Traffic.....	44
3.6	Deployment Troubleshooting.....	44
<b>4</b>	<b>Upgrading Software Version .....</b>	<b>45</b>
<b>5</b>	<b>Downgrading Software Version .....</b>	<b>47</b>
<b>6</b>	<b>Licensing Mediant CE .....</b>	<b>48</b>
6.1	Obtaining and Activating a Purchased License Key .....	48
6.2	Installing the License Key .....	49
6.3	Product Key.....	50

## List of Figures

Figure 1-1: Mediant CE Architecture .....	7
Figure 1-2: Signaling Components Switchover .....	8
Figure 1-3: Network Architecture for Mediant CE Deployment in Single Availability Zone .....	9
Figure 1-4: Network Architecture for Mediant CE Deployment in Multiple Availability Zones .....	11
Figure 1-5: Virtual IP Address in AWS Route Table .....	13
Figure 1-6: Virtual IP Address in AWS Route Table After Switchover .....	13
Figure 1-7: Virtual IP Address in Mediant VE IP Interfaces Table .....	13
Figure 1-8: Connection via Virtual IP Addresses through the AWS Transit Gateway .....	15
Figure 1-9: Virtual IP Addresses in AWS Transit Gateway Route Table .....	15
Figure 1-10: Use of AWS Network Load Balancer .....	17
Figure 2-1: Searching for Mediant VE Product in the AWS Marketplace .....	21
Figure 2-2: Mediant VE Product in AWS Marketplace .....	22
Figure 2-3: Creating Route Table .....	24
Figure 2-4: Creating Cluster Subnet .....	25
Figure 2-5: Changing Cluster Subnet Route Table .....	25
Figure 2-6: Editing Route Table Association .....	26
Figure 2-7: Creating Private EC2 Endpoint .....	27
Figure 2-8: Creating NAT Gateway .....	28
Figure 2-9: Editing Route Table .....	29
Figure 2-10: Creating Default Route .....	29
Figure 3-1: Stack Manager Main Screen .....	31
Figure 3-2: Create Stack Dialog – Step 1 .....	32
Figure 3-3: Create Stack Dialog – Step 2 .....	32
Figure 3-4: Create Stack Dialog – Step 3 – “Single Zone” Deployment Topology .....	33
Figure 3-5: Create Stack Dialog – Step 3 – “Multiple Zones” Deployment Topology .....	34
Figure 3-6: Create Stack Dialog – Step 4 .....	35
Figure 3-7: Create Stack Dialog – Step 5 .....	36
Figure 3-8: Successful Stack Creation .....	37
Table 3-9: Assignment of Default Security Groups to Components and Network Interfaces .....	41
Table 3-10: Inbound Rules for Default Security Groups .....	41
Table 3-11: Minimal Required Outbound Rules for Cluster Security Group .....	42
Figure 4-1: Upgrading Mediant CE via Stack Manager .....	45
Figure 6-1: Software License Activation Tool .....	48
Figure 6-2: Product Key in Order Confirmation E-mail .....	49
Figure 6-3: Viewing Product Key .....	50
Figure 6-4: Empty Product Key Field .....	50
Figure 6-5: Entering Product Key .....	50

## Notice

Information contained in this document is believed to be accurate and reliable at the time of publishing. However, due to ongoing product improvements and revisions, AudioCodes cannot guarantee accuracy of published material after the Date Published nor can it accept responsibility for errors or omissions. Updates to this document can be downloaded from <https://www.audiocodes.com/library/technical-documents>.

This document is subject to change without notice.

Date Published: May-06-2025

## Customer Support

Customer technical support and services are provided by AudioCodes or by an authorized AudioCodes Service Partner. For more information on how to buy technical support for AudioCodes products and for contact information, please visit our website at <https://www.audiocodes.com/services-support/maintenance-and-support>.

## Stay in the Loop with AudioCodes



## Abbreviations and Terminology

Each abbreviation, unless widely used, is spelled out in full when first used.

Abbreviation	Description
MC	Media Component
SC	Signaling Component

## Document Revision Record

LTRT	Description
11017	Initial document release for Version 7.6

## Documentation Feedback

AudioCodes continually strives to produce high quality documentation. If you have any comments (suggestions or errors) regarding this document, please fill out the Documentation Feedback form on our website at <https://online.audiocodes.com/documentation-feedback>.

# 1 Introduction

**Mediant Cloud Edition (CE) Session Border Controller (SBC)**, hereafter referred to as *Mediant CE*, is a software-based product that can be deployed in one of the following operational environments:

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud
- OpenStack
- Non-cloud virtual environments (e.g. VMware)

This document describes deployment of Mediant CE in an Amazon Web Services (AWS) environment.

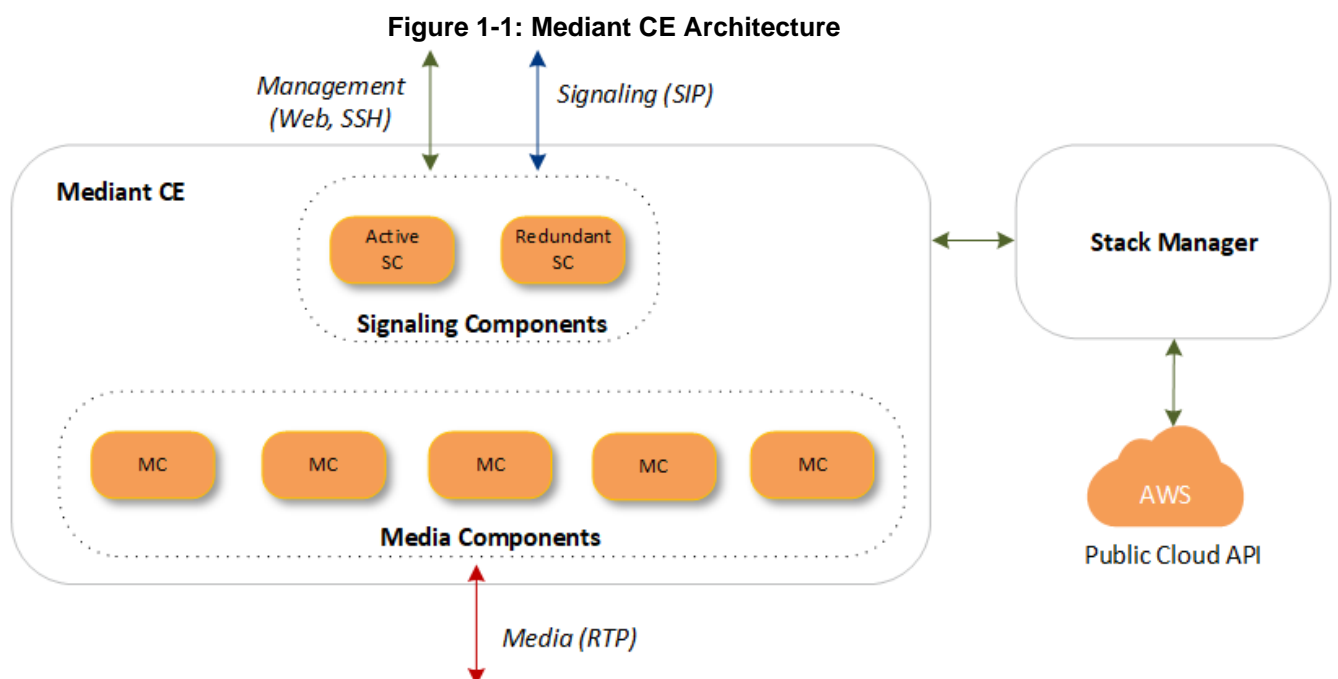
For detailed instructions on Mediant CE installation in other operational environments (for example, VMware), refer to the dedicated installation manual.



**Note:**

- The scope of this document does not fully cover security aspects for deploying the product in the AWS cloud. Security measures should be done in accordance with AWS security policies and recommendations.
- For configuring Mediant CE SBC, refer to the *Mediant Software SBC User's Manual*.

## 1.1 Architecture Overview



Mediant CE cluster is comprised of multiple components (virtual machines) that perform distinct functions:

- **Signaling Components (SC):** Handle signaling (SIP) and management (Web, SSH, etc) traffic. It also determines which MC (see below) handles the specific media traffic, which is based on load balancing between the MCs.
- **Media Components (MC):** Handle media (RTP, RTCP) traffic, including transcoding functionality. Up to 21 MCs can be used in the deployed Mediant CE.

Incoming calls are initially processed (at signaling level) by SCs, which choose the MC based on the current cluster utilization and pass the media streams to it.

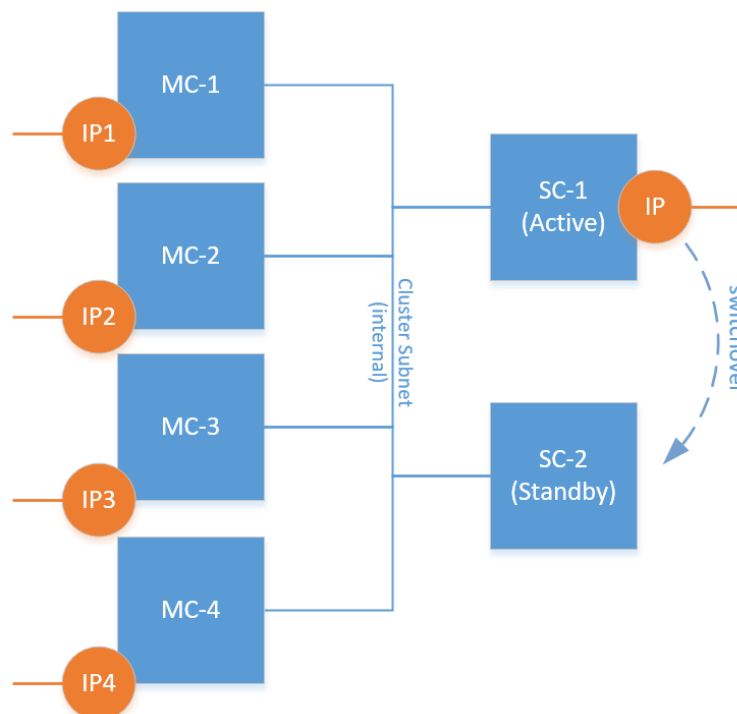
SCs also serve as a “single point of contact” for all management tasks. They provide Web and CLI interfaces through which customers have complete control over all cluster components.

## 1.2 Deployment Topology

In a typical Mediant CE deployment, two SCs are created and operate in 1+1 Active / Standby mode. In case of a failure in the active SC, all IP addresses are seamlessly moved to the remaining (newly active) SC and all established calls are preserved.

Mediant CE cluster may contain up to 21 MCs, which operate in N+1 Load Sharing mode. In case of a failure in a specific MC, calls handled by it are re-distributed across remaining MCs, with no visible effect on established calls.

**Figure 1-2: Signaling Components Switchover**



It is possible to adjust cluster size by scaling MCs “in” or “out” based on cluster utilization and/or explicit customer request. “Scaled down” MCs are kept in “shutdown” state, ensuring that they can be quickly started during “scale out” operations.





**Note:** Mediant VE and CE products share the same software image. AudioCodes has published the image for these products on AWS Marketplace under the name "Mediant VE Session Border Controller (SBC)". Therefore, in some places in this document, this product name is referenced even though the document concerns Mediant CE.

### 1.3 Stack Manager

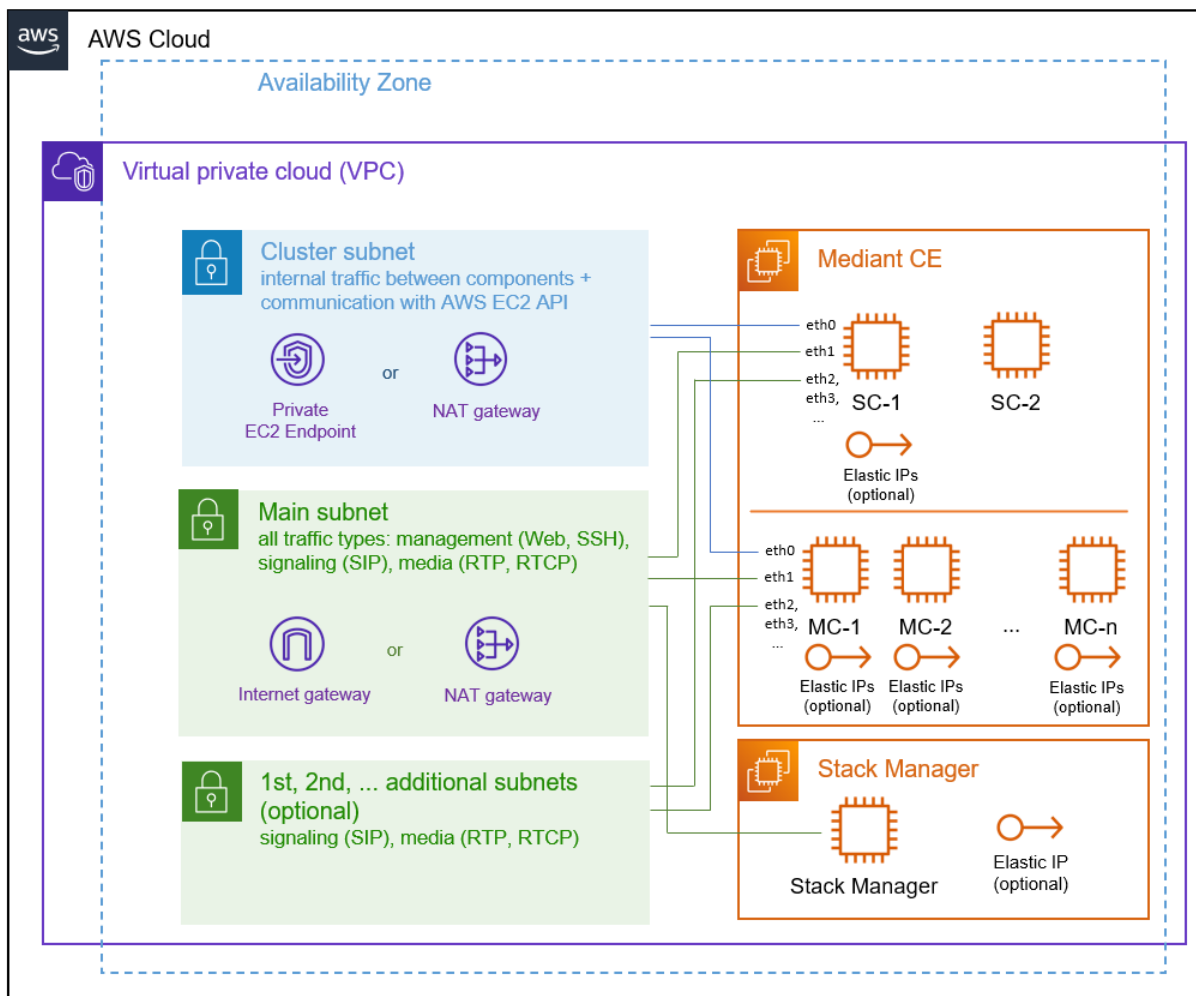
The Stack Manager tool is provided as part of the solution. It is used for initial Mediant CE cluster deployment and complete lifecycle management. For example, update of network topology, rebuild of cluster components in case of underlying cloud resources corruption or accidental removal etc.

Stack Manager also supports automatic scaling of MCs based on cluster utilization, thus significantly reducing associated infrastructure costs.

### 1.4 Deployment In a Single Availability Zone

The following diagram shows network architecture for high-available Mediant CE deployment in single Availability Zone.

**Figure 1-3: Network Architecture for Mediant CE Deployment in Single Availability Zone**



Virtual Private Cloud (VPC) must have the following subnets defined prior to Mediant CE deployment:

- **Cluster Subnet:** Carries internal communication between Mediant CE components. It is connected to both SC and MCs as the first network interface (eth0). For security reasons it is recommended to create dedicated cluster subnet and protect it from access by other instances / equipment.  
Cluster subnet is also used by SCs for accessing the AWS EC2 APIs during activity switchover. Therefore, it should have either **Private EC2 Endpoint** or **NAT Gateway** attached (for more information, see Section Cluster Subnet).
- **Main Subnet:** Carries management (HTTP, SSH, etc.), signaling (SIP) and media (RTP, RTCP) traffic; connected to both SC and MCs as the second network interface (eth1). It is also recommended to connect Stack Manager instance to the Main subnet, to enable seamless access from it to the Mediant CE's management interfaces. Since Stack Manager requires access to AWS EC2, CloudFormation and IAM APIs for its operation, Main subnet should have either **Internet Gateway** or **NAT Gateway** attached.
- **1<sup>st</sup>, 2<sup>nd</sup>, ... Additional Subnets:** Carry signaling (SIP) and media (RTP, RTCP) traffic; connected to SC and MCs as the third, fourth etc. network interfaces (eth2, eth3, ...) correspondingly; these subnets are optional, as the Main Subnet may carry all types of traffic.

All subnets must reside in the same Availability Zone of the Virtual Private Cloud (VPC).

During deployment, Stack Manager creates all relevant Mediant CE components, including SC and MC instances and Elastic IP addresses.

- **Signaling Components (SC)**

Two SCs are deployed and operate in 1+1 Active/Standby mode. Active SC uses secondary addresses on all network interfaces (except for the 1st one, connected to Cluster subnet) to communicate with external entities. During switchover, the newly active SC reassigns these addresses to itself by communicating with AWS EC2 API via the Cluster subnet.

Elastic IPs may be assigned to relevant SC network interfaces to enable communication with corresponding signaling or management entities via the public addresses. During a switchover, Elastic IPs are reassigned to a newly active instance together with the corresponding secondary IP addresses.

- **Media Components (MC)**

Up to 21 MCs may be deployed and operate in N+1 Load Sharing mode. Each MC has its own set of IP addresses and uses all available IP addresses on the network interface, including the primary address. Elastic IPs may be assigned to relevant MC network interfaces to enable communication with corresponding media entities via the public addresses.

It is recommended to deploy **Stack Manager** into the same VPC and connect it to the Main subnet, to enable seamless connectivity with the deployed Mediant CE via private IP addresses. Elastic IP may be assigned to it to allow outbound access to AWS APIs and inbound access from the Internet.

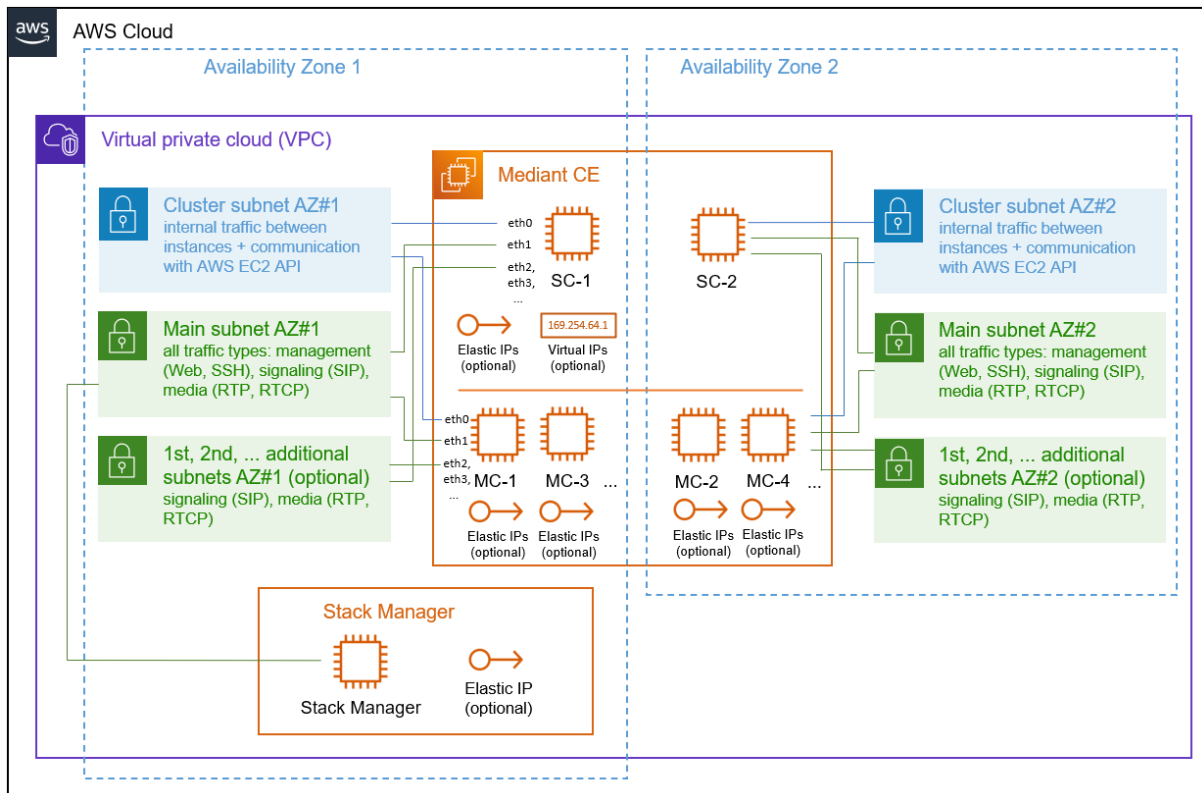
## 1.5 Deployment In Multiple Availability Zone



**Note:** Deployment in multiple availability zones is supported starting from version 7.4.500.

The following diagram shows network architecture for high-available Mediant CE deployment in multiple Availability Zones.

**Figure 1-4: Network Architecture for Mediant CE Deployment in Multiple Availability Zones**



The deployment requires two sets of subnets – Cluster, Main and optionally Additional 1, Additional 2, etc. – that must be defined in two different Availability Zones and created prior to Mediant CE deployment. Refer to the previous section for detailed description of what traffic is carried by each subnet and what are corresponding pre-requisites (e.g. Cluster subnet must have either Private EC2 Endpoint or NAT Gateway attached).

Mediant CE components are distributed across two Availability Zones – half of them resides in the Availability Zone 1 and is connected to the first set of subnets, another half resides in the Availability Zone 2 and is connected to the second set of subnets.

■ **Signaling Components (SC)**

Two SCs are deployed and operate in 1+1 Active/Standby mode. Each SC resides in a different Availability Zone and is connected to the corresponding subnets set.

Each SC uses its own set of IP addresses, including primary addresses, on each network interface. Communication with external equipment via public IP addresses (e.g. over Internet) is performed via Elastic IP addresses, that are assigned to the Active instance and reassigned to another (newly active) instance during activity switchover.

Communication with external equipment via private IP addresses (inside VPC or via Transit Gateway) is performed via Virtual IP addresses – private IP addresses outside the VPC address space, that are defined in the AWS subnet’s routing table and whose route destination is updated during activity switchover. For a detailed description, see Section 1.5.1.

Reassignment of Elastic and Virtual IP addresses is done by the newly active SC by communicating with AWS EC2 API via the Cluster subnet.

■ **Media Components (MC)**

Up to 21 MCs may be deployed and operate in N+1 Load Sharing mode. Half of the MCs reside in the first Availability Zone and another half in the second Availability Zone. Each MC has its own set of IP addresses and uses all available IP addresses on network interface, including primary address. Elastic IPs may be assigned to relevant MC network interfaces to enable communication with corresponding media entities via the public addresses.

It is recommended to deploy **Stack Manager** into the same VPC and connect it to the Main subnet in one of the Availability Zones to enable seamless connectivity with the deployed Mediant CE via private IP addresses. An elastic IP may be assigned to it to allow outbound access to AWS APIs and inbound access from the Internet.



**Note:** Mediant CE deployment in multiple availability zones can use AWS Network Load Balancer instead of Virtual and/or Elastic IP addresses. For a detailed description, see Section 1.5.2’AWS Network Load Balancer.

## 1.5.1 Virtual IP Addresses

Mediant CE deployment in multiple availability zones uses Virtual IP addresses to enable communication between deployed Mediant CE Signaling Components and other equipment via private IP addresses.

Virtual IP addresses are special IP addresses that must reside outside the VPC address space. Stack Manager allocates them by default from 169.254.64.0/24 subnet, thereby ensuring that they don’t collide with your VPC range. For production deployments, it’s recommended to allocate your own virtual IP addresses and specify them using the **virtual\_ip\_sc\_ethX** advanced configuration parameter, for example:

```
virtual_ip_sc_eth2 = 10.1.5.15
```



**Note:** If you manually specify Virtual IP addresses, make sure that they reside outside the VPC address space. For example, if your VPC CIDR is 172.31.0.0/16, you can’t specify 172.31.100.11 as the Virtual IP address because it resides within the VPC address space.

Virtual IP addresses are “manually plugged” by Stack Manager into the routing tables of subnets attached to the corresponding interfaces of deployed Signaling Components. The entry is for a specific IP address (prefix /32) and is initially configured to point to the network interface of the 1<sup>st</sup> Signaling Component, which is initially active.

**Figure 1-5: Virtual IP Address in AWS Route Table**

Destination	Target	Status	Propagated
2a05:d014:f3c:5a00::/56	local	Active	No
0.0.0.0/0	igw-0a49ae63	Active	No
20.0.0.0/24	pcx-000f02cee1024a314	Active	No
169.254.64.1/32	eni-0c36cc91ef8ffeb81	Active	No
172.31.0.0/16	local	Active	No

Upon an HA switchover, the new active Signaling Component updates the entry with its own network interface, thus ensuring that all communication via the Virtual IP address is sent towards it.

**Figure 1-6: Virtual IP Address in AWS Route Table After Switchover**

Destination	Target	Status	Propagated
2a05:d014:f3c:5a00::/56	local	Active	No
0.0.0.0/0	igw-0a49ae63	Active	No
20.0.0.0/24	pcx-000f02cee1024a314	Active	No
169.254.64.1/32	eni-02cc5638a6de644ed	Active	No
172.31.0.0/16	local	Active	No

Virtual IP addresses are defined in Mediant CE’s Signaling Components’ IP Interfaces table, and applications (e.g., SIP Interface) are “bound” to them.

**Figure 1-7: Virtual IP Address in Mediant VE IP Interfaces Table**

IP Interfaces (4)

+ New Edit | Page 1 of 1 | Show 10 records per page

INDEX	NAME	APPLICATION TYPE	INTERFACE MODE	IP ADDRESS	PREFIX LENGTH	DEFAULT GATEWAY	PRIMARY DNS	SECONDARY DNS	ETHERNET DEVICE
0	eth0	MAINTENANCE	IPv4 Manual	172.31.89.147	20	172.31.80.1	172.31.0.2	0.0.0.0	vlan 1
1	eth1	OAMP + Media + Con	IPv4 Manual	172.31.68.92	20	172.31.64.1	172.31.0.2	0.0.0.0	vlan 2
2	eth2	Media + Control	IPv4 Manual	172.31.11.83	20	172.31.0.1	172.31.0.2	0.0.0.0	vlan 3
3	eth2:200	Media + Control	IPv4 Manual	169.254.64.1	32	172.31.0.1	172.31.0.2	0.0.0.0	vlan 3

Stack Manager automatically “plugs” virtual IP addresses into the AWS route tables attached to the corresponding network interfaces of both deployed Signaling Components. If you want it to update additional AWS route tables within the same VPC, specify them using the following advanced configuration parameter:

```
additional_route_tables = eth1:rtb-123,eth2:rtb-567|rtb-890
```



**Note:** You should not specify the Transit Gateway route table in the **additional\_route\_tables** advanced configuration parameter. Instead, refer to the following section for detailed instructions.

**Known limitations:**

- Allocation of multiple Virtual IP addresses on the same network interface is not supported.

### 1.5.1.1 Using Virtual IP Addresses for Communication Across VPCs and with On-Prem Networks

The presence of Virtual IP addresses in AWS route tables attached to the corresponding subnets ensures that other equipment deployed in the same AWS subnets (or in subnets with the same route tables) can communicate with the deployed Median CE’s Signaling Components via these Virtual IP addresses (communication is within the VPC). For example, if you have an IP-PBX or a Contact Center deployed in the same AWS subnet as Mediant CE SBC, it can use Virtual IP addresses to communicate with them.


If equipment (e.g., IP-PBX or Contact Center) that needs to communicate with the Mediant CE SBC resides in a different AWS VPC or in the on-premise network, connected via AWS Direct Connect or site-to-site VPN, you must use **AWS Transit Gateway** to establish proper connectivity with Mediant VE SBC via the Virtual IP addresses.



**Note:** Regular VPC peering, routes only IP addresses belonging to the corresponding VPC CIDR ranges and therefore, doesn’t support communication via Virtual IP addresses. Use AWS Transit Gateway and configure it as described below to enable communication via Virtual IP addresses across VPCs.

AWS Transit Gateway is a network transit hub which enables connectivity between AWS VPCs and on-premise networks. Transit Gateway has a route table that can be explicitly configured to route Virtual IP addresses – a specific one or the whole subnet range – to the specific VPC network attachment.

Figure 1-8: Connection via Virtual IP Addresses through the AWS Transit Gateway



➤ To configure AWS Transit Gateway for proper connectivity via Virtual IP addresses:

1. Open the AWS VPC console (<https://console.aws.amazon.com/vpc>).
2. Navigate to the **Transit Gateways** screen, and then select your Transit Gateway.
3. Under the **Details** tab, locate the associated route table and navigate to it.
4. Under the **Routes** tab:
  - a. Click **Create static route**.
  - b. For **CIDR**, enter the subnet range from which Virtual IP addresses are allocated, for example, 169.254.64.0/24.
  - c. For **attachment**, choose the attachment that represents the VPC where Mediant VE is deployed.
  - d. Click **Create static route**.

Figure 1-9: Virtual IP Addresses in AWS Transit Gateway Route Table

CIDR	Attachment ID	Resource ID	Resource type	Route type	Route state
100.0.0.0/24	tgw-attach-0311181a3b5380f83	vpc-05fe37c58e4430716	VPC	Propagated	Active
169.254.64.0/24	tgw-attach-0d981a54f2459e027	vpc-45f3152c	VPC	Static	Active
172.31.0.0/16	tgw-attach-0d981a54f2459e027	vpc-45f3152c	VPC	Propagated	Active

You also need to manually update the route tables of the subnets where the equipment that communicates with Mediant CE resides so that they route Virtual IP addresses (a specific one or the whole range) to the AWS Transit Gateway.

For connection between VPCs, do the following:

1. In the AWS VPC console, navigate to the **Subnets** screen, and then select the subnet (in a different VPC) where other equipment resides.
2. Under the **Details** tab, locate the route table and navigate to it.
3. Under the **Routes** tab, click **Edit routes**:
  - a. Click **Add route**.
  - b. For **Destination**, enter the subnet range from which Virtual IP addresses are allocated, for example, 169.254.64.0/24.
  - c. For **Target**, choose the AWS Transit Gateway configured above.
  - d. Click **Save changes**.
4. Repeat the above steps for all applicable subnets.

Finally, you need to verify that the default route table in the VPC where Mediant VE is deployed is configured with Virtual IP addresses.

1. Open the AWS VPC console (<https://console.aws.amazon.com/vpc>).
2. Navigate to the **VPC** screen, and then select the VPC where Mediant VE is deployed.
3. Under the **Details** tab, locate the main route table and navigate to it.
4. Under the **Routes** tab, check the presence of the routes to Virtual IP addresses.

If the main route table was attached to the subnet where Virtual IP addresses were allocated, it will already contain proper routes and no additional configuration is needed. Otherwise, you need to configure Mediant VE SBC to update the main route table:

- a. Open Stack Manager.
- b. Navigate to your Mediant VE stack.
- c. Click **Modify**, and then enter the following in the advanced config section:

```
additional_route_tables = <if-name>:<route-table-id>
```

For example:

```
additional_route_tables = eth1:rtb-123
```

If you have Virtual IP addresses configured on multiple network interfaces, specify multiple entries for each network interface separated by a comma, for example:

```
additional_route_tables = eth1:rtb-123,eth2:rtb-123
```

- d. Click **Update** to apply the changes.



## 1.5.2 AWS Network Load Balancer

Mediant CE deployment in multiple availability zones can use AWS Network Load Balancer (NLB) instead of Virtual and/or Elastic IP addresses.



**Note:** This section describes an *alternative* method of Mediant CE deployment in multiple availability zones. Refer to the description below for detailed information on pros and cons of this deployment method.

The following operational modes are supported:


- NLB is not used. This is the default mode.
- Internal NLB is used instead of Virtual IP addresses. This is the recommended alternative method.
- Public NLB is used instead of Elastic IP addresses.
- Internal / Public NLBs are used instead of both Virtual and Elastic IP addresses.

AWS NLB uses DNS Name (FQDN) as it's frontend addresses.

In a typical Mediant CE deployment, NLB is comprised of two internal nodes. Each node resides in a different availability zone and has its own IP address. NLB's frontend DNS Name (FQDN) is resolved into one of these two IP addresses (of the internal nodes).

NLB maintains the status of the Signaling Component (SC) instances that reside behind it using keep-alive probes. At any time, one (active) SC instance is "alive" and the other (standby) SC instance is not. Therefore, NLB sends the received traffic to the active SC instance. When an HA switchover happens, the newly active SC instance becomes "alive" and NLB starts sending the traffic to it.

**Figure 1-10: Use of AWS Network Load Balancer**



AWS NLB is a standard AWS Networking element and therefore, is fully compatible with various networking topologies and components, including VPC peering, Direct Connect Gateways, etc.

The caveat of using the AWS NLB is that it uses DNS Name as its frontend address. This DNS Name is resolved into one of the two different IP addresses (of internal NLB nodes). This means that equipment that communicates with the Mediant CE SBC via the NLB must support the use of DNS Names / FQDNs and be able to perform discovery of new IP addresses in a timely manner. Note that NLB is used for signaling / management traffic only; media streams flow directly to the Media Components.

**Known limitations:** IPv6 addresses can only be used for TCP/TLS traffic when AWS NLB is used. This is because IPv6 addresses are implemented via “dualstack” NLB mode, where frontend DNS Name / FQDN is resolved into both A (IPv4) and AAAA (IPv6) records, while still using IPv4 addresses to communicate with the backend SC instances; and “dualstack” NLB mode doesn’t support UDP rules.

The following lists the pros and cons of using the AWS NLB as opposed to regular Elastic and Virtual IP addresses:

■ **Compatibility with various AWS networking elements / topologies:**

- Elastic IP addresses and AWS NLB are native AWS networking components and as such are fully supported by various AWS networking topologies and components (e.g., VPC peering, Direct Connect gateway, etc.).
- Virtual IP addresses are entries in AWS route tables, manually created during SBC deployment and updated during a switchover. They must reside outside the VPC address range and require AWS Transit Gateway for any traffic that flows outside the subnet. Virtual IP addresses are not compatible with VPC peering and certain VPN topologies.

■ **Switchover mechanism:**

- AWS NLB uses keep-alive messages to discover the status of SC instances and sends traffic to the active SC instance accordingly.
- Elastic IP addresses and Virtual IP addresses are relatched to the active SC instance using AWS APIs. These APIs may exhibit delays when the datacenter is overloaded.

■ **Interoperability with SIP / management equipment:**

- AWS NLB uses DNS Name as its frontend address. This DNS name is resolved into one of the two IP addresses (of internal NLB nodes). Use of NLB mandates that external SIP / management equipment that communicates with Mediant CE supports the use of DNS Name and is able to refresh the address resolution in a timely manner.
- Elastic and Virtual IP addresses are moved during a switchover. Therefore, SIP / management equipment that communicates with the SBC works with a single set of IP addresses, regardless of which SC instance is currently active.

■ **IPv6 addresses:**

- AWS NLB supports only TCP/TLS traffic in “dualstack” mode used for IPv6 addresses.
- IPv6 addresses are assigned to the active SBC instance and moved during a switchover. All types of traffic – UDP / TCP / TLS – are supported.

Use the above list to determine whether AWS NLB suits your deployment needs or not. It's recommended to consider the option of using AWS NLB instead of Virtual IP addresses only, while keeping Elastic IPs for communication over the public IP addresses.

### 1.5.2.1 Prerequisites for AWS Network Load Balancer Deployment

If you use AWS Network Load Balancer (NLB) for Mediant CE deployment, update the IAM role assigned to Stack Manager to include the following action:

```
"elasticloadbalancing:*"
```

This is needed to allow Stack Manager to create AWS NLB and all associated resources.

If you receive the following error message during creation of Mediant CE stack with AWS NLB:

```
"User is not authorized to perform iam:CreateServiceLinkedRole
on resource   arn:aws:iam::<account-id>:role/aws-service-
role/elasticloadbalancing.amazonaws.com/..."
```

create a corresponding service linked role through the AWS CLI:

```
aws iam create-service-linked-role --aws-service-name
elasticloadbalancing.amazonaws.com
```

or add the following to the IAM role assigned to Stack Manager:

```
{
  "Effect": "Allow",
  "Action": "iam:CreateServiceLinkedRole",
  "Resource": "arn:aws:iam::*:role/aws-service-
role/elasticloadbalancing.amazonaws.com/*",
  "Condition": {"StringLike": {"iam:AWSServiceName":
"elasticloadbalancing.amazonaws.com"}}
}
```

### 1.5.2.2 Enabling AWS Network Load Balancer

To enable use of AWS Network Load Balancer (NLB), use the following advanced config parameter during stack creation:

```
ha_nlb = internal - use internal NLB instead of Virtual IPs
ha_nlb = public   - use public NLB instead of Elastic IPs
ha_nbl = all      - use internal / public NLBs instead of
                  Virtual / Elastic IP addresses
```

### 1.5.2.3 Configuration for AWS Network Load Balancer Deployment

When you use AWS Network Load Balancer (NLB) for Mediant CE multi-zone deployment, equipment that communicates with the SBC via the corresponding network interfaces (e.g., SIP Contact Center, IP-PBX, or management system) must be configured to use NLB DNS Name / FQDN.

The latter can be found in the AWS dashboard or **More > Show IP Addresses** action in the Stack Manager,

You must also configure the “Local Hostname” parameter for the IP Groups that communicate via the AWS NLB to contain the NLB DNS Name / FQDN value. This adds the NLB FQDN to the Contact header of SIP messages and ensures that SIP messages traverse via the NLB and therefore, always reach the active SBC instance regardless of which SBC instance was active during call establishment.



**Note:** Failure to configure NLB DNS Name / FQDN as the “Local Hostname” for IP Groups that communicate via the AWS NLB prevents SIP sessions from properly maintaining connection after a switchover.

## 2 Installation Prerequisites

Prior to installing Mediant CE in the AWS environment, make sure that you meet the following prerequisites:

- You have an AWS account. If you don't have an AWS account, you can sign up for one on Amazon's website at <http://aws.amazon.com/>.
- You have subscribed to the AudioCodes Mediant VE offer in AWS Marketplace. For more information, see Section [Subscribing to AudioCodes Mediant VE Product in AWS Marketplace](#).
- You have created an Identity and Access Management (IAM) role that enables Mediant CE to manage its network interfaces. For more information, see Section [IAM Role for Mediant CE](#).
- You have created all subnets needed for Mediant CE deployment, including the Cluster subnet with a private EC2 endpoint or NAT gateway. For more information, see Section [Cluster Subnet](#).

### 2.1 Subscribing to AudioCodes Mediant VE Product in AWS Marketplace

Mediant VE and CE products share the same software image. AudioCodes distributes Mediant VE/CE software images by publishing them in the AWS Marketplace.

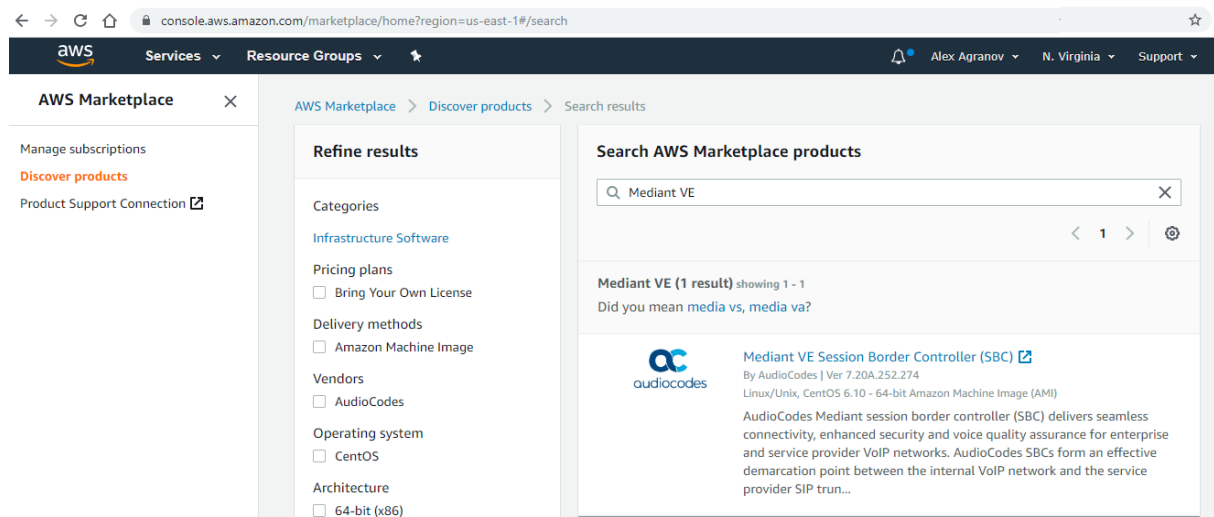


**Note:** As Mediant VE and CE products share the same software image, AudioCodes has published the image for these products on AWS Marketplace under the name "Mediant VE Session Border Controller (SBC)".

Prior to deploying the Mediant CE you must subscribe to the AudioCodes Mediant VE product in AWS Marketplace as follows:

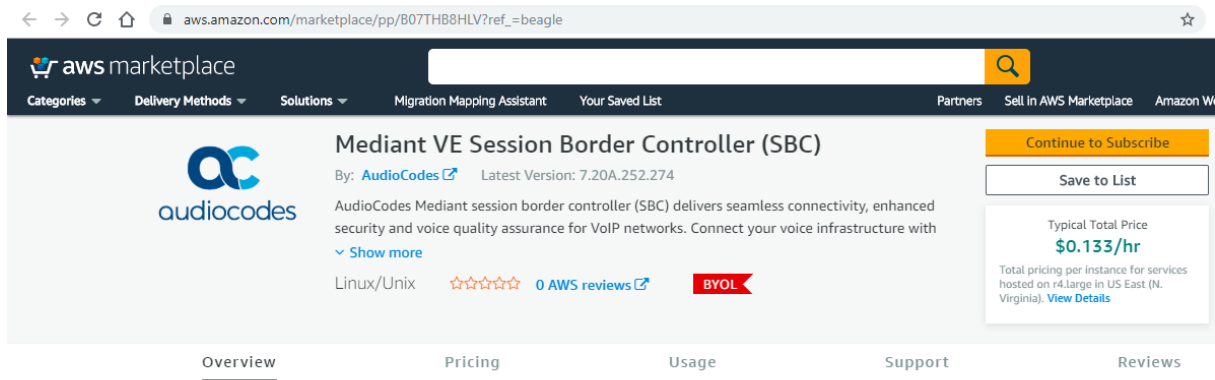
1. Open the AWS Marketplace console at <https://console.aws.amazon.com/marketplace>.
2. In the **Discover Products** tab, search for the "Mediant VE" product.

**Figure 2-1: Searching for Mediant VE Product in the AWS Marketplace**



3. Click the **Mediant VE Session Border Controller (SBC)** product.

Figure 2-2: Mediant VE Product in AWS Marketplace



### Product Overview

AudioCodes Mediant session border controller (SBC) delivers seamless connectivity, enhanced security and voice quality assurance for enterprise and service provider VoIP networks.

AudioCodes SBCs form an effective demarcation point between the internal VoIP network and the service provider SIP trunk, performing SIP and WebRTC signaling mediation, translation and media handling (better known as interoperability), while also securing your VoIP solution.

AudioCodes SBCs can connect virtually any existing VoIP infrastructure and IP-PBX to Amazon Chime Voice Connector, Microsoft Teams or Skype for Business environments, enabling coexistence and simple migration to cloud-based solutions.

#### Highlights

- Easily secure your VoIP environment and connect to any SIP provider
- Tested to work with Amazon Chime Voice Connector
- Certified for Microsoft Teams Direct Routing and Skype for Business

4. Click **Continue to Subscribe** to subscribe to the Mediant VE product.

## 2.2 IAM Role for Mediant CE

The following IAM role must be created prior to creating the Mediant CE stack. This role ensures that Mediant CE components can manage their network interfaces and re-assign IP addresses in case of a switchover.

The role differs depending on HA deployment topology – “single zone” or “multiple zones”. Note that “multiple zones” deployment topology is supported starting from Version 7.4.500.

### ➤ IAM Role for “single zone” Mediant CE deployment:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:AssignPrivateIpAddresses",
        "ec2:UnassignPrivateIpAddresses",
        "ec2:AssociateAddress",
        "ec2:DescribeAddresses",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeNetworkInterfaces"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

➤ **IAM Role for “multiple zones” Mediant CE deployment:**

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:AssociateAddress",
        "ec2:DescribeAddresses",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeNetworkInterfaces",
        "ec2:ReplaceRoute"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

➤ **To create an IAM Role:**

1. Open the AWS IAM management console at <https://console.aws.amazon.com/iam>.
2. Navigate to the **Policies** screen:
  - a. Click **Create**.
  - b. Select the **JSON** tab, copy-and-paste the IAM policy rules listed above, and then click **Review policy**.
  - c. Enter the IAM policy name (e.g. "SBC\_HA"), and then click **Create policy**.
3. Navigate to the **Roles** screen:
  - a. Click **Create role**.
  - b. Choose **EC2** use case, and then click **Next: permissions**.
  - c. Search for the IAM policy created in the previous step, select it, and then click **Next: tags**.
  - d. Click **Next: review**.
  - e. Enter the IAM role name (e.g. "SBC\_HA"), and then click **Create role**.

## 2.2.1 IAM Role for Initial Configuration from S3 URL

Mediant CE SBC may be provided with an initial configuration INI file, stored on AWS Simple Storage Service (S3), during its launch. This is done by specifying `sc_ini_file_url` advanced configuration parameter during stack deployment.

If you use this option, add the following rules to the IAM Role created previously, to enable Mediant CE SBC access to the corresponding S3 bucket (replace “sbc” in the example below with the actual bucket name).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
    }
  ],
}
```

```

        "Resource": "arn:aws:s3:::sbc"
    },
    {
        "Effect": "Allow",
        "Action": [
            "s3:GetObject",
        ],
        "Resource": "arn:aws:s3:::sbc/*"
    }
]
}

```

## 2.3 Cluster Subnet

The Cluster Subnet is used for the following tasks:

- Internal communication between Mediant CE components
- Accessing AWS API (for IP address management)

Mediant CE uses private addresses in the Cluster Subnet. Therefore, to enable Mediant CE to access AWS API via the Cluster subnet, you must do one of the following:

- Create a private EC2 endpoint in the Cluster subnet (recommended method)
- Attach a NAT gateway to the Cluster subnet (alternative method)

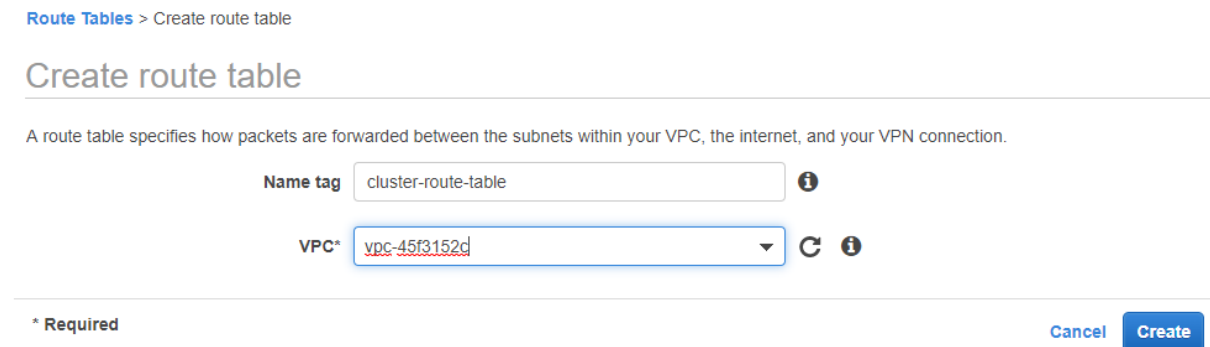
In addition, since the Cluster subnet carries sensitive information, it's recommended to create a dedicated subnet and protect it from unauthorized access.

### 2.3.1 Creating Cluster Subnet

➤ **To create the Cluster subnet:**

1. Open the AWS VPC management console at <https://console.aws.amazon.com/vpc>.
2. Open the Route Tables page, and then click **Create route table**:
  - a. In the 'Name tag' field, enter the new route table name (e.g. 'cluster-route-table').
  - b. From the 'VPC' drop-down list, select the VPC where Mediant CE will be deployed.
  - c. Click **Create** to create the route table.

**Figure 2-3: Creating Route Table**





3. Open the Subnets page, and then click **Create Subnet**.
  - a. In the 'Name tag' field, enter the new subnet name (e.g. 'cluster-subnet').
  - b. From the 'Availability Zone' drop-down list, select the Availability Zone where Mediant CE will be deployed.
  - c. In the 'IPv4 CIDR block' field, enter the IPv4 CIDR for the subnet.
  - d. Click **Yes, Create** to create the route table.

**Figure 2-4: Creating Cluster Subnet**

Subnets > Create subnet

### Create subnet

Specify your subnet's IP address block in CIDR format; for example, 10.0.0.0/24. IPv4 block sizes must be between a /16 netmask and /28 netmask, and can be the same size as your VPC. An IPv6 CIDR block must be a /64 CIDR block.

Name tag  ⓘ

VPC\*  ⓘ

VPC CIDRs

CIDR	Status	Status Reason
172.31.0.0/16	associated	

Availability Zone  ⓘ

IPv4 CIDR block\*  ⓘ

\* Required Cancel **Create**

4. Select the created subnet, switch to the **Route Table** tab, and then click **Edit route table association**.

**Figure 2-5: Changing Cluster Subnet Route Table**

Subnet: subnet-035888fc2f2e95bf8 ☰ ☰ ☰

Description  Flow Logs  **Route Table**  Network ACL  Tags  Sharing

**Edit route table association**

Route Table: rtb-379b7d5e

⏪ < 1 to 2 of 2 > ⏩


Destination	Target
172.31.0.0/16	local
0.0.0.0/0	<a href="#">igw-0a49ae63</a>

5. Choose the Cluster route table created in the previous steps, and then click **Save**.

Figure 2-6: Editing Route Table Association

## Edit route table association

Subnet ID subnet-0496039603680f5a2

Route Table ID\*  

1 to 2 of 2	
Destination	Target
172.31.0.0/16	local

\* Required Cancel



**Note:** Make sure that Cluster subnet has a dedicated route table. Other subnets (Main subnet, Additional subnets) should be attached to different route table(s), which would typically have the Internet Gateway configured as the default route to ensure proper functionality of Elastic IPs attached to the corresponding network interfaces of EC2 instances.

### 2.3.2 Creating Private EC2 Endpoint in Cluster Subnet

After you have successfully created the Cluster subnet, you need to enable access to the AWS API via through this subnet. The recommended method is to create a private EC2 endpoint in the Cluster subnet.

- **To create the private EC2 endpoint in Cluster subnet:**
  1. Open the **Security Groups** page, and then click **Create security group**.
    - a. In the 'Security group name' field, enter the security group name (e.g., "Endpoint Security Group").
    - b. In the 'VPC' drop-down list, select the VPC where Mediant CE will be deployed.
    - c. Under 'Inbound rules', click **Add rule** and then configure the rule as follows:
      - ◆ 'Type': Custom TCP
      - ◆ 'Port range': 443
      - ◆ 'Source': Anywhere
    - d. Click **Create security group** to create the new security group.
  2. Open the **Endpoints** page, and then click **Create Endpoint**.
    - a. In the 'Service Category' field, select **AWS services**.
    - b. In the 'Service Name' field, enter "ec2" in the search box and then press Enter. Select the EC2 endpoint from the list (e.g., **com.amazonaws.eu-central-1.ec2**).
    - c. In the 'VPC' drop-down list, select the VPC where Mediant CE will be deployed.
    - d. In the 'Subnets' field, select the HA subnet.
    - e. Select the 'Enable DNS name' checkbox.

- f. In the 'Security group' field, remove the default security group and then select the 'Endpoint Security Group' that you created in the previous step.
- g. Click **Create Endpoint** to create the new endpoint.

**Figure 2-7: Creating Private EC2 Endpoint**

Endpoints > Create Endpoint

### Create Endpoint

A VPC endpoint allows you to securely connect your VPC to another service.  
 An interface endpoint is powered by [PrivateLink](#), and uses an elastic network interface (ENI) as an entry point for traffic destined to the service.  
 A gateway endpoint serves as a target for a route in your route table for traffic destined for the service.

- Service category**
- AWS services
  - Find service by name
  - Your AWS Marketplace services

**Service Name** com.amazonaws.eu-central-1.ec2 ⓘ

⌕ < 1 to 50 of more >

Service Name	Owner	Type
<input type="radio"/> com.amazonaws.eu-central-1.codebuild	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.codecommit	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.codepipeline	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.config	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.datasync	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.dynamodb	amazon	Gateway
<input checked="" type="radio"/> com.amazonaws.eu-central-1.ec2	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.ecr.api	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.transfer.server	amazon	Interface
<input type="radio"/> com.amazonaws.eu-central-1.workspaces	amazon	Interface

**VPC\*** vpc-45f3152c ⓘ

**Subnets** subnet-0496039603680f5a2 ⓘ

Availability Zone	Subnet ID
<input type="checkbox"/> eu-central-1a (eu1-az2)	subnet-78c72611
<input checked="" type="checkbox"/> eu-central-1b (eu1-az3)	subnet-0496039603680f5a2 (cluster)
<input type="checkbox"/> eu-central-1c (eu1-az1)	subnet-42be9e08

**Enable DNS name**  Enable for this endpoint ⓘ

To use private DNS names, ensure that the attributes 'Enable DNS hostnames' and 'Enable DNS Support' are set to 'true' for your VPC (vpc-45f3152c). [Learn more.](#)

**Security group** sg-8a7791e3 ⓘ [Create a new security group](#)

\* Required

Cancel **Create endpoint**

### 2.3.3 Creating NAT Gateway in Cluster Subnet



**Note:** If you created a Private EC2 Endpoint in the Cluster subnet, as described in the previous section, you can skip this section.

An alternative method for enabling access to the AWS API through the Cluster subnet is by attaching a NAT Gateway to the Cluster subnet.

➤ **To create NAT Gateway and attach it to the Cluster subnet:**

1. Open the NAT Gateways page, and then click **Create NAT Gateway**:
  - a. From the 'Subnet' drop-down list, select a subnet that belongs to the same Availability Zone where the Cluster subnet was created (and where Mediant CE will be deployed) and that has an Internet Gateway attached to it. For example, select **Main Subnet**.



**Note:** Do not select **Cluster Subnet** at this stage. The NAT Gateway itself will be configured as a default route in the Cluster Subnet and therefore, it won't be able to access the Internet from it.

- b. From the 'Elastic IP Allocation ID' drop-down list, select an existing Elastic IP if you have pre-allocated Elastic IPs in your VPC, or click **Create New EIP** to create a new one.
- c. Click **Create a NAT Gateway** to create the NAT gateway.

**Figure 2-8: Creating NAT Gateway**

[NAT Gateways](#) > Create NAT Gateway

#### Create NAT Gateway

Create a NAT gateway and assign it an Elastic IP address. [Learn more.](#)

Subnet\* subnet-be6e8bc3

Elastic IP Allocation ID\* eipalloc-067ef98ad76079011

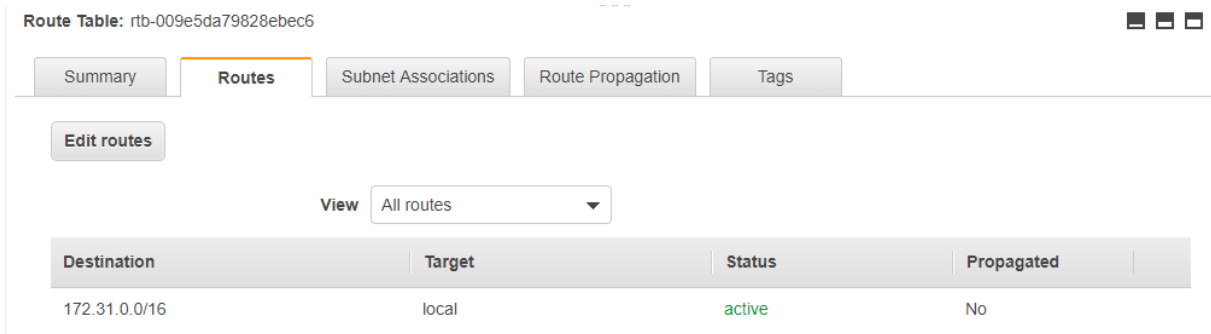
New EIP (3.122.83.211) creation successful.



\* Required

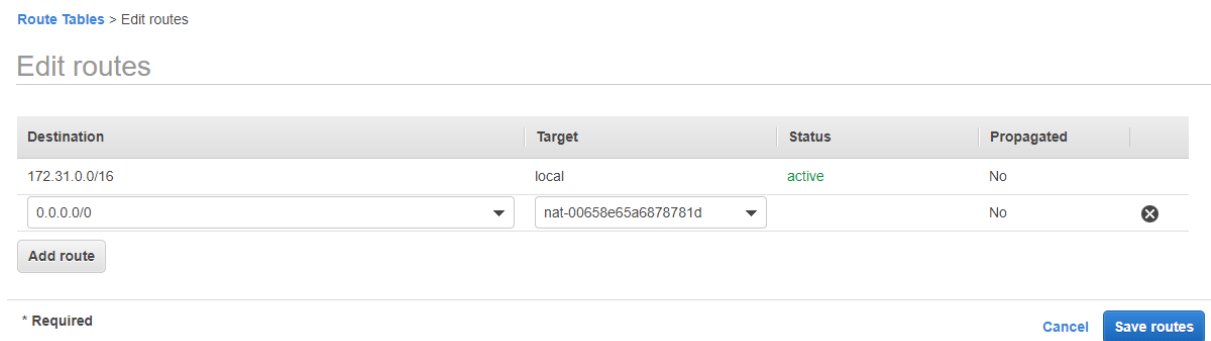
2. Open the Route Tables page, and then select the Cluster route table created in the previous steps.
3. Switch to the **Routes** tab, and then click **Edit routes** to edit the routes.

**Figure 2-9: Editing Route Table**



4. Create the default route entry (0.0.0.0/0) that points to the created NAT gateway, and then click **Save** to save your changes.

**Figure 2-10: Creating Default Route**



## 2.4 Instance Types

The default Mediant CE deployment uses the following instance types:

- **SC instances:** m5.2xlarge
- **Forwarding MC instances:** m5n.large (for three network interfaces or less) or m5n.xlarge (for four network interfaces)
- **Transcoding MC instances:** c5.4xlarge

You may customize instance types during stack creation.

Refer to the *Release Notes* for a complete list of instance types supported by Mediant CE, their capacities and capabilities.

## 3 Deploying Mediant CE

Deployment of Mediant CE on AWS is performed via the Stack Manager.

Stack Manager is a management tool developed by AudioCodes that enables simple and intuitive deployment and complete lifecycle management of Mediant VE and Mediant CE products on public clouds. The tool provides the following features for Mediant CE:

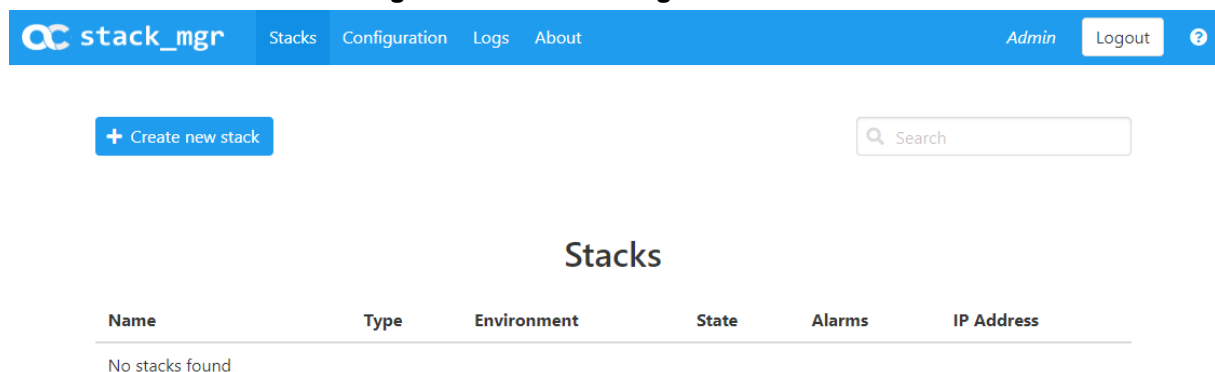
- Initial product deployment.
- Update of deployed stack's network topology
- Automatic and on-demand scaling of MCs to adjust stack footprint and minimize infrastructure costs.
- Monitoring of deployed AWS resources and recovery in case of their corruption / accidental removal
- Upgrade of software on all Mediant CE components
- Removal of all deployed resources in case of stack deletion

Stack Manager uses dynamically generated Cloud Formation templates for stack deployment on AWS platform and is not involved in call processing or any other service provided by the Mediant CE.

### ➤ To deploy Mediant CE:

1. Install the Stack Manager tool, as described in the *Stack Manager User's Manual*.
2. Log into the Stack Manager tool after deployment; the following screen appears:

**Figure 3-1: Stack Manager Main Screen**



3. Click **Create** to create a new stack; the following dialog box appears:

Figure 3-2: Create Stack Dialog – Step 1

4. In the 'Name' field, enter the name of the stack (e.g., “mediant-ce”).
5. From the 'Stack type' drop-down list, select **Mediant CE**.
6. From the 'Region' drop-down list, select the region where the stack will be deployed; additional fields appear:

Figure 3-3: Create Stack Dialog – Step 2

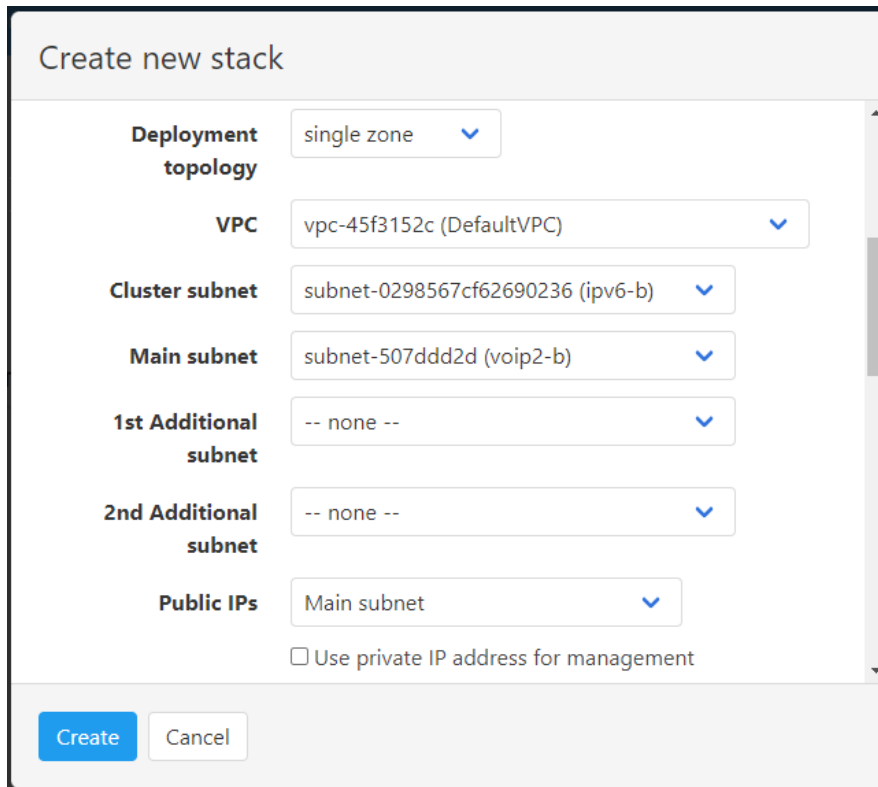
7. From the 'Key pair' drop-down list, select the key pair to access the deployed stack's CLI interface (via SSH protocol), or leave it at the default value (**none**) if you plan to use username / password (defined in the same dialog later) to access both Web and CLI interfaces.
8. From the 'IAM role' drop-down list, select the IAM role that corresponds to the deployment topology (**single zone** or **multiple zones**). See Section IAM Role for Mediant CE for details.



9. For HA deployment, from the 'Deployment topology' drop-down list, select **single zone** or **multiple zones**.

From the 'VPC' drop-down list, select the VPC where the stack will be deployed; additional fields appear:

**Figure 3-4: Create Stack Dialog – Step 3 – “Single Zone” Deployment Topology**



The screenshot shows a dialog box titled "Create new stack" with the following fields and options:

- Deployment topology:** single zone (dropdown)
- VPC:** vpc-45f3152c (DefaultVPC) (dropdown)
- Cluster subnet:** subnet-0298567cf62690236 (ipv6-b) (dropdown)
- Main subnet:** subnet-507ddd2d (voip2-b) (dropdown)
- 1st Additional subnet:** -- none -- (dropdown)
- 2nd Additional subnet:** -- none -- (dropdown)
- Public IPs:** Main subnet (dropdown)
- Use private IP address for management

At the bottom, there are two buttons: "Create" (blue) and "Cancel" (white).

10. Select the subnets that Mediant CE will be connected to. The list of subnets depends on the deployment topology that you selected before. For example, if you selected **multiple zones** deployment topology, two sets of subnets appears for each Availability Zone.

Figure 3-5: Create Stack Dialog – Step 3 – “Multiple Zones” Deployment Topology

The screenshot shows a 'Create new stack' dialog box with the following configuration:

- Deployment topology:** multiple zones (selected)
- VPC:** vpc-45f3152c (DefaultVPC)
- Cluster subnet:**
  - (zone 1): -- select --
  - (zone 2): -- select --
- Main subnet:**
  - (zone 1): -- select --
  - (zone 2): -- select --
- 1st Additional subnet:**
  - (zone 1): -- select --

Buttons: Create, Cancel

11. For **Public Ips**, select which subnets need to communicate with external equipment via public IP addresses. Stack Manager assigns Elastic IP addresses to the corresponding network interfaces.
12. If you assign a Public IP address to the Main subnet, Stack Manager by default uses this public IP address for communicating with the deployed stack. You may override this behaviour by checking the 'Use private IP address for management' checkbox. In this case, Stack Manager uses the private IP address to communicate with the deployed stack.

Figure 3-6: Create Stack Dialog – Step 4

Create new stack

Signaling Components

VM type t3.large  Customize

Media Components

Profile forwarding

VM type t3.small  Customize

Min number 2

Max number 3

Admin User

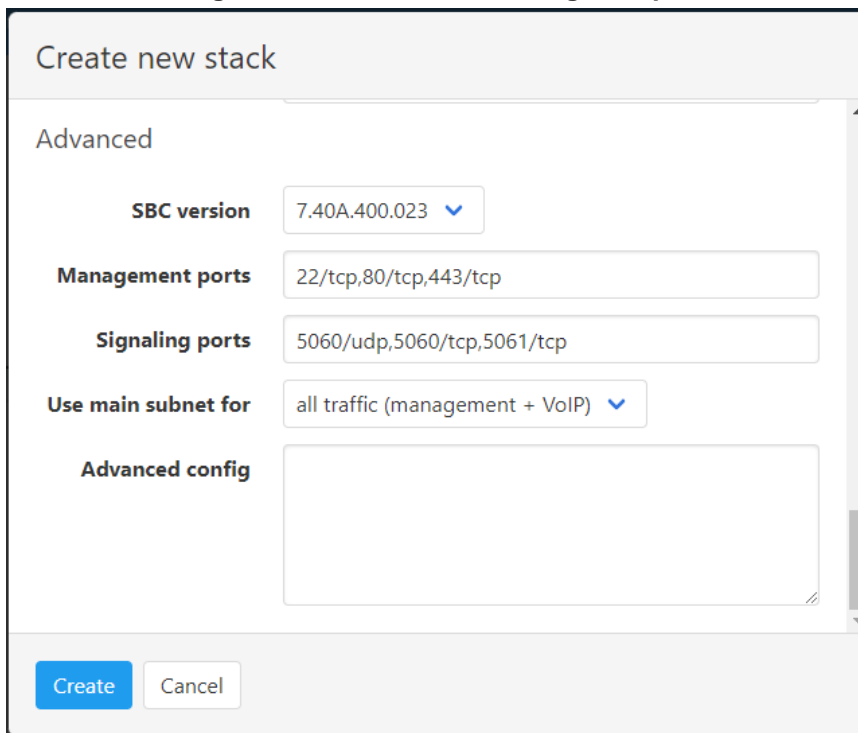
Username

Password

Create Cancel

13. The **VM type** for both SCs and MCs is pre-selected and automatically updated based on other parameters that you defined in the Create Stack dialog box. If you want to modify it, select the 'Customize' checkbox next to it and select a value from the drop-down list.
14. From the 'Profile' drop-down list, select whether you need MCs to perform simple media stream **forwarding** (includes RTP-to-SRTP translation and vice versa) or **transcoding** capabilities (for coder conversion or DTMF detection).
15. From the 'Min number' and 'Max number' drop-down lists, select the minimum and maximum number of MCs in the stack. Stack Manager creates a maximum number of MCs, but initially starts only a minimum number of them. You may later adjust the number of running MCs via **scale out** and **scale in** actions.
16. In the 'Username' and 'Password' fields, enter the admin user credentials that will be configured on the deployed stack. You use these credentials when connecting to the stack via Web or CLI management interfaces. Note that Stack Manager uses different credentials to communicate with the stack – **StackMgr** user and randomly generated password. Therefore, even if you later change admin user credentials (e.g., via Mediant CE's Web or CLI interface) communication between Stack Manager and the deployed Mediant CE stack will not be affected.

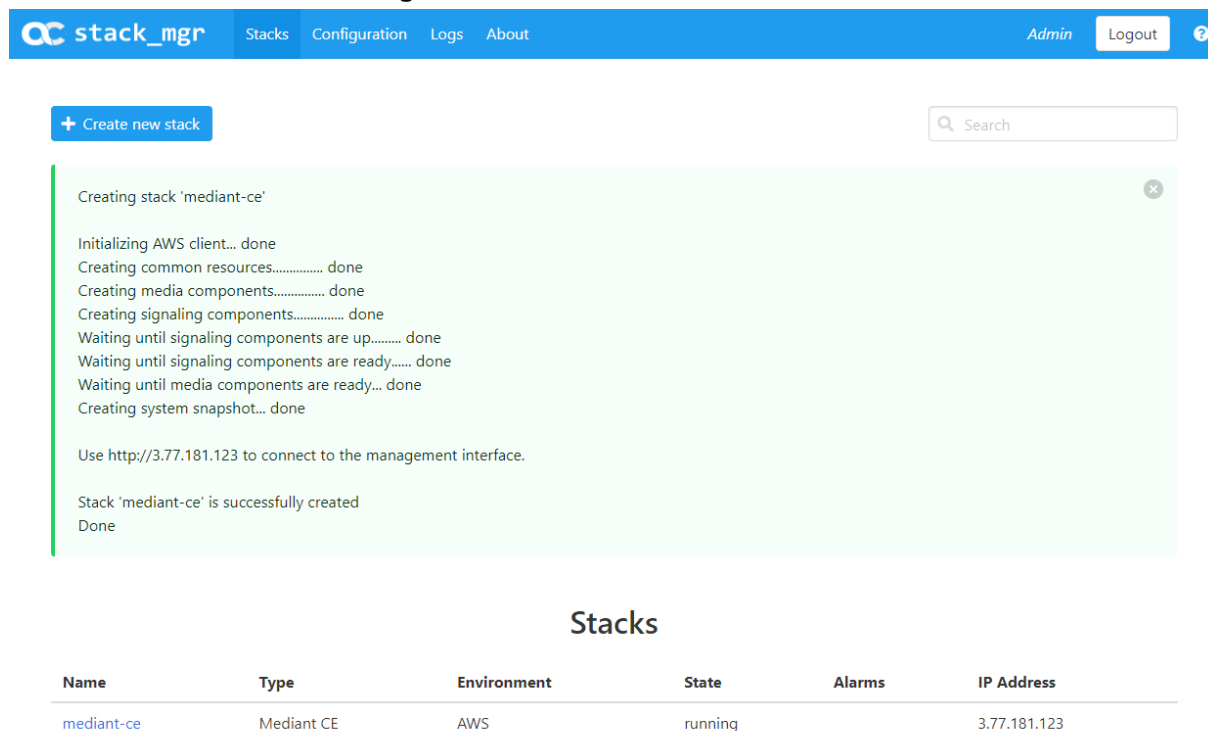
Figure 3-7: Create Stack Dialog – Step 5



17. From the 'SBC version' drop-down list, select the Mediant CE version that you want to deploy.
18. In the 'Management ports' and 'Signaling ports' fields, enter the list of management and signaling ports respectively that should be open on the Mediant CE. Specified ports are configured in the corresponding network security groups assigned to Mediant CE network interfaces. The value is a comma-separated list of the following elements:
  - <port>/udp: Opens specific UDP port for all sources (e.g., 161/udp)
  - <port>/udp/<cidr>: Opens a specific UDP port for traffic originating from specific CIDR (e.g., 161/udp/172.16.0.0/16 opens UDP port 161 for traffic from 172.16.0.0/16 subnet)
  - <port>/tcp: Opens a specific TCP port (e.g., 22/tcp)
  - <port>/tcp/<cidr>: Opens a specific TCP port for traffic originating from specific CIDR (e.g., 22/tcp/172.16.1.0/24)
  - icmp: Opens ICMP traffic
19. From the 'Use main subnet for' drop-down list, choose whether the Main subnet should be used for all traffic (management, signaling and media) or for management traffic only. This effects network security groups assigned to the Mediant CE network interface connected to the Main subnet and default SIP Interface and Media Realm created on Mediant CE.
20. In the 'Advanced config' text box, enter advanced configuration parameters, if needed. See the next sections for a partial list of the supported advanced configuration parameters. Refer to the *Stack Manager User's Manual* for a complete list.

21. Click **Create** to start stack creation.
22. Wait until the stack is created.

Figure 3-8: Successful Stack Creation



The screenshot shows the Stack Manager web interface. At the top, there is a navigation bar with 'stack\_mgr', 'Stacks', 'Configuration', 'Logs', and 'About' menus, along with 'Admin' and 'Logout' buttons. Below the navigation bar, there is a '+ Create new stack' button and a search box. The main content area displays a green box with the following text:

```

Creating stack 'mediant-ce'

Initializing AWS client... done
Creating common resources..... done
Creating media components..... done
Creating signaling components..... done
Waiting until signaling components are up..... done
Waiting until signaling components are ready..... done
Waiting until media components are ready... done
Creating system snapshot... done

Use http://3.77.181.123 to connect to the management interface.

Stack 'mediant-ce' is successfully created
Done
  
```

Below the green box, there is a table titled 'Stacks' with the following columns: Name, Type, Environment, State, Alarms, and IP Address. The table contains one entry:

Name	Type	Environment	State	Alarms	IP Address
mediant-ce	Mediant CE	AWS	running		3.77.181.123

## 3.1 Public IP Addresses

During Mediant CE stack creation, Stack Manager lets you specify which subnets (and corresponding network interfaces) will be assigned with public (Elastic) IP addresses via the **Public IPs** parameter in the **Networking** section.

Stack Manager by default applies the same configuration for both SCs and MCs. For each assigned public (Elastic) IP address, it creates corresponding entries in the NAT Translation configuration table (of Mediant CE), thus ensuring that when the SIP application attached to the corresponding private IP addresses communicates with external SIP peers, it essentially does this via the Elastic IP address.

It is possible to specify different configuration for SCs and MCs. It is also possible to attach multiple Elastic IP addresses to the same network interface. This may be done by specifying the **sc\_public\_ips** and/or **mc\_public\_ips** advanced configuration parameter in the **Advanced Config** section during stack creation, or updating **SC Public IPs** and/or **MC Public IPs** parameters for existing stack via the **Modify** action.



**Note:** When the **sc\_public\_ips** or **mc\_public\_ips** advanced configuration parameter is specified in the **Advanced Config** section during stack creation, it overrides any value configured via the **Public IPs** parameter in the **Networking** section for the corresponding components (SC or MC).

### ■ **sc\_public\_ips**

Contains a comma-separated list of subnet names (main, additional1, additional2, etc.), which will be assigned with public (Elastic) IP addresses, and optionally, the number of public (Elastic) IP addresses on the corresponding network interface.

For example:

```
sc_public_ips = main:2,additional1
```

attaches two Elastic IP addresses to the network interface connected to the Main subnet (eth1) and one Elastic IP address to the network interface connected to the Additional 1 subnet (eth2).

■ **mc\_public\_ips**

Same as above, but for MC network interfaces.

For example:

```
mc_public_ips = main,additional1:2
```

Stack Manager automatically creates secondary private IP addresses on the network interfaces that may be required for public (Elastic) IP attachment. The exact behavior depends on the component type and deployment topology:

■ **SCs:**

- For deployments in a **single availability zone**, Elastic IP addresses are always attached to the secondary private IP addresses. For each Elastic IP address, corresponding secondary IP addresses are implicitly created.
- For deployments in **multiple availability zones**, the first Elastic IP address is attached to the primary private IP address. For each additional Elastic IP address, corresponding secondary IP addresses are implicitly created.

- **MCs:** First Elastic IP address is attached to the primary private IP address. For each additional Elastic IP address, corresponding secondary IP addresses are implicitly created.

## 3.2 Private IP Addresses for Deployments in Single Availability Zone



**Note:** This section applies only to deployments in a **single availability zone** (as described in Section 1.4).

Stack Manager always creates one "operational" private IP address on each network interface. The exact behavior depends on the component type:

- **SCs:** Primary IP addresses on eth1, eth2 and eth3 interfaces (connected to Main, 1<sup>st</sup> and 2<sup>nd</sup> Additional subnets correspondingly) are not used, because they can't be moved between two SC instances during activity switchover; instead, secondary IP addresses are created and used
- **MCs:** Primary IP addresses are used on each interface

For subnets that have an assigned public (Elastic) IP address (as described in Section 3.1), the Elastic IP address is mapped (by AWS) to the corresponding interface's first "operational" private IP address, and the NAT Translation table is configured (on Mediant CE) to reflect this mapping.

For subnets that don't have an assigned public (Elastic) IP address, communication with other equipment (inside the VPC or between "connected" VPCs) happens via the interface's first "operational" private IP address.

If you want to enable communication via both public (Elastic) and private IP addresses on the same subnet or add multiple private IP addresses to a network interface, specify **sc\_additional\_ips** and/or **mc\_additional\_ips** advanced configuration parameters in the

**Advanced Config** section during stack creation, or update **SC Additional IPs** and/or **MC Additional IPs** parameters for existing stack via the **Modify** action.

■ **sc\_additional\_ips**

Contains a comma-separated list of subnet names (main, additional1, additional2, etc.), which will be assigned with additional private IP addresses, and optionally, the number of additional private IP addresses on the corresponding network interface.

For example:

```
sc_additional_ips = main,additional1:2
```

attaches one additional private IP address to the network interface connected to the Main subnet (eth1) and two additional private IP addresses to the network interface connected to the Additional 1 subnet (eth2).

■ **mc\_additional\_ips**

Same as above, but for MC network interfaces.

For example:

```
mc_additional_ips = main,additional1:2
```

The number of additional private IP addresses specified via the **sc\_additional\_ips** or **mc\_additional\_ips** advanced configuration parameter is added *on top* of any private IP address created by Stack Manager by default and/or due to the public (Elastic) IP addresses assigned to the specific network interface.

For example, the following configuration:

```
Deployment Topology: single zone
Cluster Subnet: <cluster-subnet-id>
Main Subnet: <main-subnet-id>
1st Additional Subnet: <additional-subnet-id>
Public IPs: Main subnet
Advanced Config:
    sc_additional_ips = main,additional1
```

creates the following networking configuration on SCs:

- **eth0** – one primary IP addresses (used for internal communication between SC instances and for communication with AWS EC2 API) and one secondary IP address (used for internal communication with MC instances)
- **eth1** – one primary and two secondary IP addresses:
  - primary IP address is not used because it can't be moved between SC instances in case of switchover
  - 1<sup>st</sup> secondary IP address – first "operational" private IP address, created implicitly and assigned with Elastic IP address (due to the **Public IPs** configuration parameter)
  - 2<sup>nd</sup> secondary IP address – created due to the **sc\_additional\_ips** advanced configuration parameter containing "main" element
- **eth2** – one primary and two secondary IP addresses:
  - primary IP address is not used because it can't be moved between SC instances in case of switchover
  - 1<sup>st</sup> secondary IP address – first "operational" private IP address, created implicitly
  - 2<sup>nd</sup> secondary IP address – created due to the **sc\_additional\_ips** advanced configuration parameter containing "additional1" element

### 3.3 Private IP Addresses for Deployments in Multiple Availability Zones



**Note:** This section applies only to deployments in **multiple availability zones** (as described in Section 1.5).

Deployments in multiple availability zones use primary private IP addresses for both signaling and media components on all network interfaces.

For subnets that have an assigned public (Elastic) IP address (as described in Section 3.1), the Elastic IP address is mapped (by AWS) to the corresponding interface's primary private IP address, and the NAT Translation table is configured (on Mediant CE) to reflect this mapping.

For subnets that don't have an assigned public (Elastic) IP address, communication with other equipment (inside the VPC or between "connected" VPCs) happens as follows:

- **For SCs:** Virtual IP addresses (see Section 1.5.1) are allocated and should be used for communication inside the VPC or between VPCs connected via Transit Gateway. Private IP addresses may appear in the IP Interfaces table, but shouldn't be used by applications (e.g., SIP Interfaces).
- **For MCs:** Private IP addresses are used for communication inside the VPC or between "connected" VPCs.

If you want to enable communication via both public (Elastic) and private IP addresses on the same subnet or add multiple private IP addresses to a network interface on Media Components, specify the **sc\_additional\_ips** and/or **mc\_additional\_ips** advanced configuration parameters in the **Advanced Config** section during stack creation, or update the **SC Additional IPs** and/or **MC Additional IPs** parameters for the existing stack via the **Modify** action.

- **sc\_additional\_ips**

Contains a comma-separated list of subnet names (main, additional1, additional2, etc.), for which virtual IP address are allocated in addition to the public (Elastic) IP address(es).

For example, the following configuration attaches both public (Elastic) and virtual IP address to the Main subnet (eth1):

```
Deployment Topology: multi zone
Public IPs: Main subnet
Advanced Config:
    sc_additional_ips = main
```

- **mc\_additional\_ips**

Contains a comma-separated list of subnet names (main, additional1, additional2, etc.), which are assigned with additional private IP addresses, and optionally, the number of additional private IP addresses on the corresponding network interface.

The specified number of additional private IP addresses is added *on top* of any private IP addresses created by Stack Manager by default, and/or due to the public (Elastic) IP addresses assigned to the specific network interface.

For example, the following configuration attaches one additional private IP address to the network interface connected to the Main subnet (eth1) and two additional private IP addresses to the network interface connected to the Additional 1 subnet (eth2):

```
mc_additional_ips = main,additional1:2
```



## 3.4 Security Groups

### 3.4.1 Default Security Groups

Stack Manager creates the following security groups during Mediant CE deployment:

- **OAM** – Security group for management traffic
- **Signaling** – Security group for SIP traffic
- **Media** – Security group for RTP/RTCP traffic
- **Cluster** – Security group for internal traffic between Mediant CE instances

These default security groups are assigned to the following components and network interfaces:

**Table 3-9: Assignment of Default Security Groups to Components and Network Interfaces**

Security Group	Subnet Names	Components	Interface Name
<b>Cluster</b>	Cluster	Signaling Components, Media Components	eth0
<b>OAM</b>	Main	Signaling Components	eth1
<b>Signaling</b>	Main, 1 <sup>st</sup> Additional, 2 <sup>nd</sup> Additional, ...	Signaling Components	eth1, eth2, eth3, ...
<b>Media</b>	Main, 1 <sup>st</sup> Additional, 2 <sup>nd</sup> Additional, ...	Media Components	eth1, eth2, eth3, ...

The following inbound rules are created for the default security groups:

**Table 3-10: Inbound Rules for Default Security Groups**

Security Group	Traffic	Protocol	Port	Source
<b>OAM</b>	SSH	TCP	22	0.0.0.0/0
	HTTP	TCP	80	0.0.0.0/0
	HTTPS	TCP	443	0.0.0.0/0
<b>Signaling</b>	SIP over UDP	UDP	5060	0.0.0.0/0
	SIP over TCP	TCP	5060	0.0.0.0/0
	SIP over TLS	TCP	5061	0.0.0.0/0
<b>Media</b>	RTP, RTCP	UDP	6000-65535	0.0.0.0/0
<b>Cluster</b>	Internal	UDP	669	<b>Cluster</b> security group
	Internal	UDP	680	<b>Cluster</b> security group
	Internal	TCP	80	<b>Cluster</b> security group
	Internal	TCP	2442	<b>Cluster</b> security group

Security Group	Traffic	Protocol	Port	Source
	Internal	TCP	2424	Cluster security group
	Internal	UDP	3900	Cluster security group
	Internal	UDP	925	Cluster security group

Outbound rules are configured by default to allow all traffic.

### 3.4.2 Adjusting Default Security Groups

Default **OAM**, **Signaling** and **Media** security groups are configured by default to accept traffic from all sources, which constitutes a significant security risk. It's highly recommended to modify them after Mediant VE creation to allow inbound traffic only from specific IP addresses and/or subnets, especially for management traffic.

Note that inbound rules for the **Cluster** Security Group allow only traffic that originates from instances that are attached to this Security Group. Therefore, typically there is no need to modify them.

Such modification can be done via the following stack configuration parameters:

- **Management Ports**

Defines a list of inbound management ports and corresponding transport protocols as provided by the **OAM** Security Group.

- The value is a comma-separated list of the following elements:

```
<port>/<protocol>/ [<cidr>]
```

Where:

- <port> is an individual port number (e.g., 22) or a port range (e.g., 22-23)
- <protocol> is tcp or udp
- <cidr> is optional and can be an IP address (e.g., 10.1.2.3) or a CIDR (e.g., 10.1.0.0/16)

For example:

```
22/tcp/10.11.2.0/24,80/tcp/10.11.2.34,443/tcp
```

- **Signaling Ports**

Defines a list of inbound signaling ports and corresponding transport protocols as provided by the **Signaling** Security Group.

- **Media Ports**

Defines a list of inbound media ports and corresponding transport protocols as provided by the **Media** Security Group.

If you update outbound rules for the **Cluster** Security Group, make sure to include the following minimal required rules:

**Table 3-11: Minimal Required Outbound Rules for Cluster Security Group**

Type	Protocol	Port Range	Destination	Description
All	All	All	Cluster Security Group	Internal traffic between Mediant CE instances
HTTP	TCP	80	169.254.169.254/32	Communication with EC2 instance metadata service

Type	Protocol	Port Range	Destination	Description
HTTPS	TCP	443	A.B.C.D/32	Communication with EC2 API endpoint. Replace A.B.C.D with the IP address of the private EC2 endpoint in the HA subnet. If you use a NAT Gateway to access the public EC2 endpoint, replace the destination with 0.0.0.0/0.

### 3.4.3 Using Custom Security Groups

Instead of modifying rules of the default Security Groups created by Stack Manager, you use custom Security Groups, for example, created by your IT department.

Such configuration can be done via the following stack advanced configuration parameters:

- **cluster\_nsg\_id**

Defines a custom Security Group to be used instead of the default **Cluster** Security Group.

For example:

```
cluster_nsg_id = sg-123456
```

- **oam\_nsg\_id**

Defines a custom Security Group to be used instead of the default **OAM** Security Group.

- **signaling\_nsg\_id**

Defines a custom Security Group to be used instead of the default **Signaling** Security Group.

- **media\_nsg\_id**

Defines a custom Security Group to be used instead of the default **Media** Security Group.

Alternatively, you can assign custom network Security Groups to a specific interface of specific components via the following stack advanced configuration parameters:

- **nsg\_id\_sc\_ethX**

Defines a custom Security Group for a specific network interface on Signaling Components instead of default Security Groups. Multiple Security Groups can be specified via a comma-separated list.

For example:

```
nsg_id_sc_eth0 = sg-123456
nsg_id_sc_eth1 = sg-34567,sg-56789
```

- **nsg\_id\_mc\_ethX**

Defines a custom Security Group for a specific network interface on Media Components instead of default Security Groups. Multiple Security Groups can be specified via a comma-separated list.

For example:

```
nsg_id_mc_eth0 = sg-123456
nsg_id_mc_eth1 = sg-34567,sg-56789
```

### 3.5 Management Traffic

Mediant CE management is performed through the Web, CLI, and REST management interfaces provided by the active SC. These management interfaces are by default accessible via private or public IP addresses on the second network interface (eth1) of the active SC.

All Mediant CE management operations are performed through this management interface. There is no need to access management interfaces on other components (e.g., on MCs) and such access is blocked by default.

### 3.6 Deployment Troubleshooting

Stack Manager uses dynamically generated Cloud Formation templates to perform deployment on AWS platform.

If Mediant CE deployment fails and the error description provided by Stack Manager is not detailed enough, refer to the Cloud Formation service's detailed logs for additional information.

## 4 Upgrading Software Version



### IMPORTANT NOTICE

For upgrading Mediant CE SBC to a version using a digitally signed .cmp file, you **must** follow the upgrade prerequisites and instructions in the document [Mediant SW-90xx SBC Signed-CMP Upgrade Procedure Configuration Note](#).

You may upgrade the software version of the deployed Mediant CE using the Software Version file (.cmp) through one of the following means:

- Using Mediant CE Web interface:
  - Upgrade SCs using the Software Upgrade Wizard (**Action > Software Upgrade**).
  - Upgrade "active" (currently running) MCs using the Cluster Management page (**SETUP > IP NETWORK > MEDIA CLUSTER > Cluster Management**).
  - Upgrade "idle" (currently stopped) MCs using Stack Manager (**Update Idle MCs**).
- Using Stack Manager's Web interface:
  - Upgrade all components at once using the **Upgrade** operation

**Figure 4-1: Upgrading Mediant CE via Stack Manager**



**Note:** Make sure that the SCs have the same or later version than the MCs.

Upgrade to the 7.6 stream (7.60A.xxx.yyy versions) can be performed only from a 7.4 stream (7.40A.xxx.yyy versions).



**Note:** If you have an earlier version installed (e.g., from 7.2 stream), first upgrade to a 7.4 stream. Only afterwards, upgrade to a 7.6 stream. Refer to 7.4 version documentation for detailed upgrade instructions.

B

100%



## 5 Downgrading Software Version

The procedure for downgrading Mediant CE software version is similar to the upgrading procedure, as described in the previous section, but in the reverse order:

- You first need to downgrade the MCs.
- Afterwards, you need to downgrade the SCs

This sequence ensures that the SCs always have the same or later version than the MCs.

When downgrading from version 7.40A.100.\* or later to version 7.40A.005.\*, the following additional configuration steps must be performed prior to the downgrade:

1. Connect to the Mediant CE's CLI interface (provided by SCs) through an SSH client or a serial console.
2. Log in as an administrative user.
3. Run the following commands:

```
enable
  <password> (e.g. "Admin")
configure system
  voice-config
  TpncpEncryptionEnable = 0
  exit
exit
```

4. Reboot the SCs using the `reload now` CLI command or the Web interface's **Reset** button.
5. Wait until the MCs are connected. Verify that their displayed status is "Connected" and not "Connected (TLS)".



**Note:** The above procedure is required because the communication protocol between the SCs and MCs was changed in version 7.40A.100.\*. Failure to perform this procedure will prevent the MCs from connecting to the SCs after the latter are downgraded to the 7.40A.005.\* version.

## 6 Licensing Mediant CE

Once you have successfully installed Mediant CE, you need to obtain, activate and then install the License Key.



**Note:** Licensing is applicable only to SCs; MCs do not require licensing.

### 6.1 Obtaining and Activating a Purchased License Key

For Mediant CE to provide you with all the required capacity and features, you need to obtain and activate a License Key which enables these capabilities.



**Note:**

- License activation is intended **only** for first-time software activation upon product purchase (or if your License Key is "lost", due to whatever reason). For subsequent software feature upgrades, the License Key file is e-mailed to you after your Purchase Order has been processed.
- For Mediant CE with two SC instances, each SC instance has its own Serial Number, Product Key and License Key. Therefore, the instructions in this section must be done per SC instance.

➤ **To obtain and activate the License Key:**

1. Open AudioCodes Web-based Software License Activation tool at <https://www.audiocodes.com/swactivation>:

**Figure 6-1: Software License Activation Tool**

Home > Software License Activation

### Software License Activation

Please enter your Product Key received from AudioCodes and the fingerprint (e.g. Serial Number or Server Machine ID) that was generated as a result of your installation.  
For technical assistance, please contact AudioCodes support at [support@audiocodes.com](mailto:support@audiocodes.com).  
\*Supports CloudBond 365 version 7.2 and above.

Product Key \*

Fingerprint \*

For instructions on how to locate your product's fingerprint, please read the documentation relevant to your product

Email \*


I'm not a robot

**SUBMIT**



2. Enter the following information:
  - **Product Key:** The Product Key identifies your specific Mediant CE purchase for the purpose of subsequent communication with AudioCodes (for example, for support and software upgrades). The Product Key is provided in the Order Confirmation e-mail sent to you by AudioCodes upon your purchase, as shown in the example below:

**Figure 6-2: Product Key in Order Confirmation E-mail**



**Note:** For Mediant CE orders with two SC instances, you are provided with two Product Keys, one for each SC instance. In such cases, you need to perform license activation twice to obtain License Keys for both SC instances.

- **Fingerprint:** The fingerprint is the Mediant CE's Serial Number. The Serial Number uniquely identifies the software installation. The Serial Number is displayed in the 'Serial Number' field on the Device Information page (**Monitor** menu > **Monitor** menu > **Summary** tab > **Device Information**).
  - **Email:** Provide one or more e-mail addresses to where you want the License Key to be sent.
3. Click **Submit** to send your license activation request.
  4. Once AudioCodes processes and completes your license activation, you will receive an e-mail notification with the License Key file attached. Open the file with any text-based program (such as Notepad) and make sure that the serial number ("**S/N**") in the License Key is correct and reflects the Serial Number of your SC instance.



**Warning:** Do not modify the contents of the License Key file.

## 6.2 Installing the License Key

For installing the License Key on Mediant CE, refer to the *Mediant Software SBC User's Manual*.



**Note:** The License Key file for Mediant CE with two SC instances must contain two License Keys - one for the active SC instance and one for the redundant SC instance. Each License Key has a different serial number ("**S/N**"), which reflects the serial number of each SC instance.


## 6.3 Product Key

The Product Key identifies a specific purchase of your Mediant CE deployment for the purpose of subsequent communication with AudioCodes (e.g., for support and software upgrades). The Product Key is provided in the order-confirmation email sent to you upon your product purchase and is used for activating your license through AudioCodes Software License Activation tool.

The Product Key is included in the License Key. Once the License Key is installed, you can view the Product Key in the following Web pages:

- License Key page (**Setup** menu > **Administration** tab > **Maintenance** folder > **License Key**). The Product Key is displayed in the read-only 'Product Key' field, as shown in the example below:

**Figure 6-3: Viewing Product Key**




- Device Information page.

If your License Key was purchased in an earlier version (for example, 7.0), the 'Product Key' field may appear empty. In such a scenario, request the Product Key from your AudioCodes sales representative. Once received, do the following:


1. Open the License Key page.
2. Locate the Product Key group:

**Figure 6-4: Empty Product Key Field**



3. Click "empty"; the following appears:

**Figure 6-5: Entering Product Key**



4. In the field, enter the Product Key, and then click **Submit** (or **Cancel** to discard your entry).



**International Headquarters**

Naimi Park  
Ofra Haza 6  
Or Yehuda, Israel  
Tel: +972-3-976-4000  
Fax: +972-3-976-4040

**AudioCodes Inc.**

80 Kingsbridge Rd  
Piscataway, NJ 08854, USA  
Tel: +1-732-469-0880  
Fax: +1-732-469-2298

Contact us: <https://www.audiocodes.com/corporate/offices-worldwide>

Website: <https://www.audiocodes.com>

©2025 AudioCodes Ltd. All rights reserved. AudioCodes, AC, HD VoIP, HD VoIP Sounds Better, IPmedia, Mediant, MediaPack, What's Inside Matters, OSN, SmartTAP, User Management Pack, VMAS, VoIPerfect, VoIPerfectHD, Your Gateway To VoIP, 3GX, AudioCodes One Voice, AudioCodes Meeting Insights, and AudioCodes Room Experience are trademarks or registered trademarks of AudioCodes Limited. All other products or trademarks are property of their respective owners. Product specifications are subject to change without notice.

Document #: LTRT-11017

