

Mediant Virtual Edition (VE) SBC

Deployment in Amazon AWS

Version 7.4

Table of Contents

1	Introduction	9
2	Deployment Topologies and Methods.....	11
2.1	Standalone Deployment	12
2.2	HA Deployment in a Single Availability Zone	13
2.3	HA Deployment in Multiple Availability Zone.....	15
2.3.1	Virtual IP Addresses	16
2.3.1.1	Using Virtual IP Addresses for Communication Across VPCs and With On-Prem Networks	18
2.3.2	AWS Network Load Balancer	21
2.3.2.1	Prerequisites for AWS Network Load Balancer Deployment	23
2.3.2.2	Enabling AWS Network Load Balancer.....	23
2.3.2.3	Configuration for AWS Network Load Balancer Deployment.....	23
2.3.2.4	Mediant VE Behavior During Switchover	24
3	Prerequisites	25
3.1	Subscribing to AudioCodes Mediant VE Product in AWS Marketplace	25
3.2	CloudFormation Template for Mediant VE HA Deployment	26
3.3	IAM Role for Mediant VE HA Deployment	27
3.3.1	IAM Role for Initial Configuration from S3 URL	28
3.4	Network Architecture and Prerequisites	29
3.4.1	HA Subnet	29
3.4.1.1	Creating HA Subnet	29
3.4.1.2	Creating Private EC2 Endpoint in HA Subnet.....	31
3.4.1.3	Creating NAT Gateway in HA Subnet	33
3.5	Instance Type.....	34
3.6	Automatic Configuration	35
4	Deploying Standalone Mediant VE via AWS EC2 Console	37
4.1	Assigning Elastic IP Addresses to the Instance	45
4.2	Post-Installation Configuration of Mediant VE SBC – PAYG Product.....	47
5	Deploying High-Availability (HA) Mediant VE via CloudFormation Service	49
5.1	Deleting HA Mediant VE Deployment	54
6	Deploying Mediant VE via Stack Manager	55
6.1	Public IP Addresses	61
6.2	Private IP Addresses for Standalone Deployments	62
6.3	Private IP Addresses for HA Deployments in Single Availability Zone	63
6.4	Private IP Addresses for HA Deployments in Multiple Availability Zones	65
6.5	Deployment Troubleshooting.....	66
7	Security Groups	67
7.1	Default Security Groups	67
7.2	Adjusting Default Security Groups.....	68
7.3	Using Custom Security Groups	70
8	Upgrading the Software Version.....	71
8.1	Method 1 – Side-By-Side Deployment of New Version	72
8.2	Method 2 – Rebuild Existing Mediant VE Instance from New Image.....	73

8.2.1	Rebuilding Existing Standalone Mediant VE Instance Deployed via AWS EC2 Console from New Image	73
8.2.2	Rebuilding Existing High-Availability (HA) Mediant VE Deployed via AWS EC2 Console from New Image	76
8.2.3	Rebuilding Existing Mediant VE Deployed via Stack Manager	78
9	Licensing the Product.....	79
9.1	Obtaining and Activating a Purchased License Key	79
9.2	Installing the License Key	81
9.3	Product Key.....	81

List of Figures

Figure 2-1:	Network Architecture for Standalone Deployment.....	12
Figure 2-2:	Network Architecture for HA Deployment in a Single Availability Zone	13
Figure 2-3:	Network Architecture for HA Deployment in a Single Availability Zone	15
Figure 2-4:	Virtual IP Address in AWS Route Table	16
Figure 2-5:	Virtual IP Address in AWS Route Table After the Switchover	17
Figure 2-6:	Virtual IP Address in Mediant VE IP Interfaces Table	17
Figure 2-7:	Connection via Virtual IP Addresses across VPCs through AWS Transit Gateway	18
Figure 2-8:	Virtual IP Addresses in AWS Transit Gateway Route Table	19
Figure 2-9:	Use of AWS Network Load Balancer	21
Figure 3-1:	Searching for Mediant VE Product in the AWS Marketplace.....	25
Figure 3-2:	Creating Route Table.....	30
Figure 3-3:	Creating Cluster Subnet.....	30
Figure 3-4:	Changing Cluster Subnet Route Table	30
Figure 3-5:	Editing Route Table Association	31
Figure 3-6:	Creating Private EC2 Endpoint	32
Figure 3-7:	Creating NAT Gateway	33
Figure 3-8:	Editing Route Table	34
Figure 3-9:	Creating Default Route	34
Figure 4-1:	Searching for Mediant VE Product in the AWS Marketplace.....	37
Figure 4-2:	Mediant VE Product Page in AWS Marketplace.....	38
Figure 4-3:	Mediant VE Configuration Page in AWS Marketplace.....	38
Figure 4-4:	Mediant VE Launch Page in AWS Marketplace	39
Figure 4-5:	Choose Instance Type Page.....	40
Figure 4-6:	Configure Instance Page	40
Figure 4-7:	Add Storage Page.....	42
Figure 4-8:	Tag Instance Page	42
Figure 4-9:	Configure Security Group Page.....	43
Figure 4-10:	Review Page	44
Figure 4-11:	Elastic IPs Page	45
Figure 4-12:	Allocated IP Address.....	46
Figure 4-13:	Associate Address Window	46
Figure 4-14:	Metering License Page	47
Figure 5-1:	CloudFormation Console	49
Figure 5-2:	CloudFormation – Create Stack Page	50
Figure 5-3:	CloudFormation - Specify Details Page (Stack Name).....	50
Figure 5-4:	CloudFormation – Stack Creation Progress	52
Figure 5-5:	CloudFormation – Stack Outputs.....	52
Figure 6-1:	Stack Manager Main Screen.....	55
Figure 6-2:	Stack Manager: Create Stack Dialog – Step 1	56
Figure 6-3:	Stack Manager: Create Stack Dialog – Step 2	56
Figure 6-4:	Stack Manager: Create Stack Dialog – Step 3 – “Single Zone” HA	57
Figure 6-5:	Stack Manager: Create Stack Dialog – Step 3 – “Multiple Zones” HA	58
Figure 6-6:	Stack Manager: Create Stack Dialog – Step 4	59
Figure 6-7:	Stack Manager: Successful Stack Creation.....	60
Table 7-1:	Assignment of Default Security Groups	67

Table 7-2: Inbound Rules for Default Security Groups	67
Table 7-3: Minimal Required Outbound Rules for HA Security Group	69
Figure 8-1: Opening Web Interface's Software Upgrade Wizard	71
Figure 8-2: Finding Network Instances associated with EC2 Instance	74
Figure 8-3: Changing Termination Behavior of Network Interface	74
Figure 8-4: Choosing Existing Network Interfaces during EC2 Instance Creation	75
Figure 8-5: Updating Cloud Formation stack	77
Figure 8-6: Upgrading Mediant VE to New Image Based on OS Version 8 via Stack Manager	78
Figure 9-1: Software License Activation Tool	80
Figure 9-2: Product Key in Order Confirmation E-mail	80
Figure 9-3: Viewing Product Key	81
Figure 9-4: Empty Product Key Field	82
Figure 9-5: Entering Product Key	82

This page is intentionally left blank.

Notice

Information contained in this document is believed to be accurate and reliable at the time of printing. However, due to ongoing product improvements and revisions, AudioCodes cannot guarantee accuracy of 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: April-12-2026

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.

Related Documentation

Manual Name
Mediant Software SBC User's Manual
SBC Series Release Notes

Document Revision Record

LTRT	Description
10867	Initial document release for Version 7.4.
10871	Creating private EC2 updated.
10874	CloudFormation template download site; instance type c5.9xlarge and c4.8xlarge; new sections for post-installation configuration and troubleshooting of Mediant VE SBC – PAYG product; miscellaneous
10877	Note added to Deploying Standalone Mediant VE via AWS EC2 Console re adding/removing network interfaces or secondary IP addresses.
10911	Multi-zone deployment topology

LTRT	Description
10917	Virtual IP addresses for HA deployment in multiple availability zones.
11000	Security groups
11004	IPv6 supported
11008	New section on virtual IP addresses for communication across VPCs and with on-prem networks
11011	AWS Network Load Balancer as alternative for multiple availability zone.
11035	Command <code>debug cloud-license toggle-cl-debug</code> replaced by <code>debug license aws-payg toggle-cl-debug</code> .

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

This document describes deployment of AudioCodes' Mediant Virtual Edition (VE) Session Border Controller (SBC), hereafter referred to as *Mediant VE*, in an Amazon Web Services (AWS) environment.

For detailed instructions on Mediant VE installation in other virtual environments (for example, VMware), refer to the *Mediant VE SBC 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 VE SBC, refer to the *Mediant Software SBC User's Manual*.

This page is intentionally left blank.

2 Deployment Topologies and Methods

Mediant VE SBC is available in AWS Marketplace as two different products:

- **Mediant VE Session Border Controller (SBC):** This product includes a trial license (limited to three SBC sessions) and requires a purchase of production license from AudioCodes.
- **Mediant VE Session Border Controller (SBC) – PAYG:** This product includes a pay-as-you-go license that enables Customers to use the SBC as much as needed and pay for the actual service consumed via their AWS account billing.

Mediant VE SBC supports the following deployment topologies:

- **Standalone topology:** Mediant VE SBC is deployed on a single EC2 instance. Deployment is performed using the AWS EC2 console, as described in Section Deploying Standalone Mediant VE via AWS EC2 Console.
- **High-availability (HA) topology:** Mediant VE SBC is deployed on two EC2 instances, operating in 1+1 Active/Standby mode.
The following HA topologies are supported:
 - **"Single zone":** All components are deployed into a single Availability Zone.
 - **"Multiple zones":** Components are spread across two different Availability Zones (supported starting from Version 7.4.500).

See the following sections for detailed descriptions of each deployment topology.

Mediant VE SBC supports the following deployment methods:

- **Via AWS EC2 Console:** Applicable to standalone deployment topology only. See Section Deploying Standalone Mediant VE via AWS EC2 Console for details.
- **Via AWS Cloud Formation Template:** Applicable to single-zone HA deployment topology only. See Section Deploying High-Availability (HA) Mediant VE via CloudFormation Service for details.
- **Via Stack Manager:** Applicable to all deployment topologies. See Section Deploying Mediant VE via Stack Manager for details.

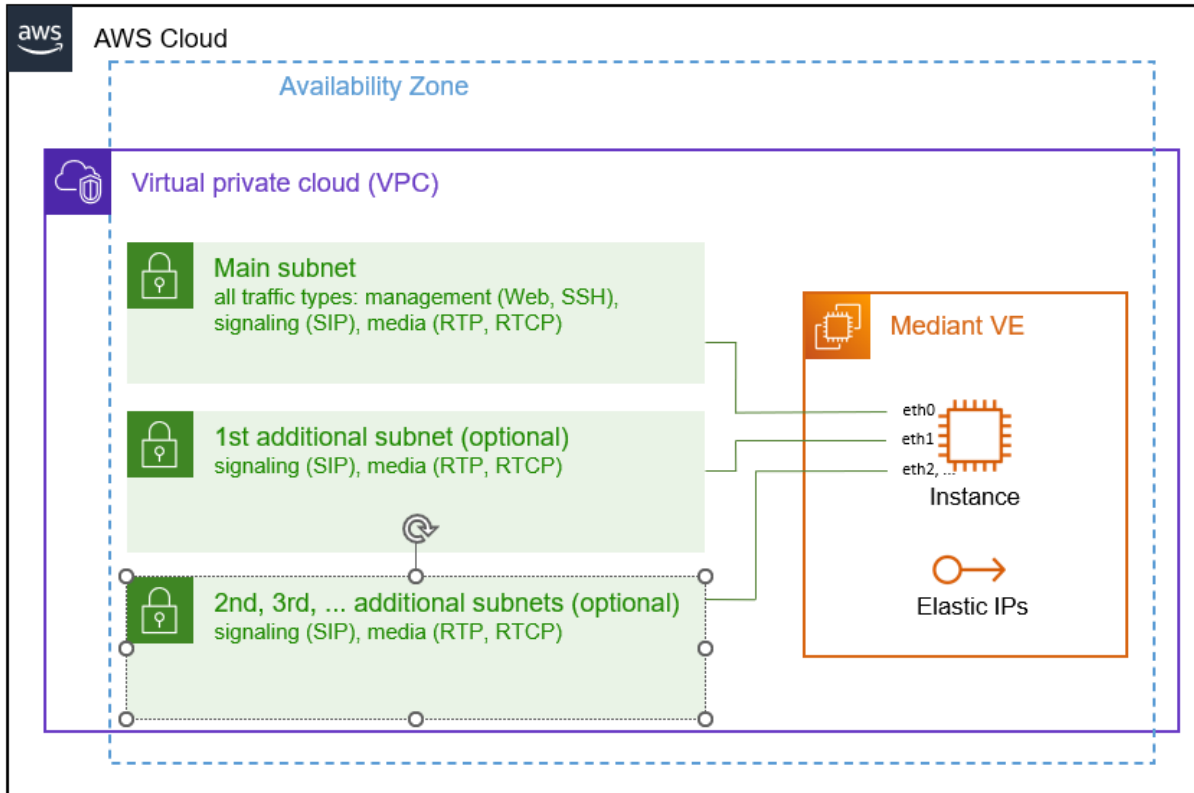


Notes: The **Mediant VE SBC – PAYG** product supports only Standalone deployment topology (not HA).

2.1 Standalone Deployment

The following diagram shows network architecture for standalone Mediant VE deployment.

Figure 2-1: Network Architecture for Standalone Deployment

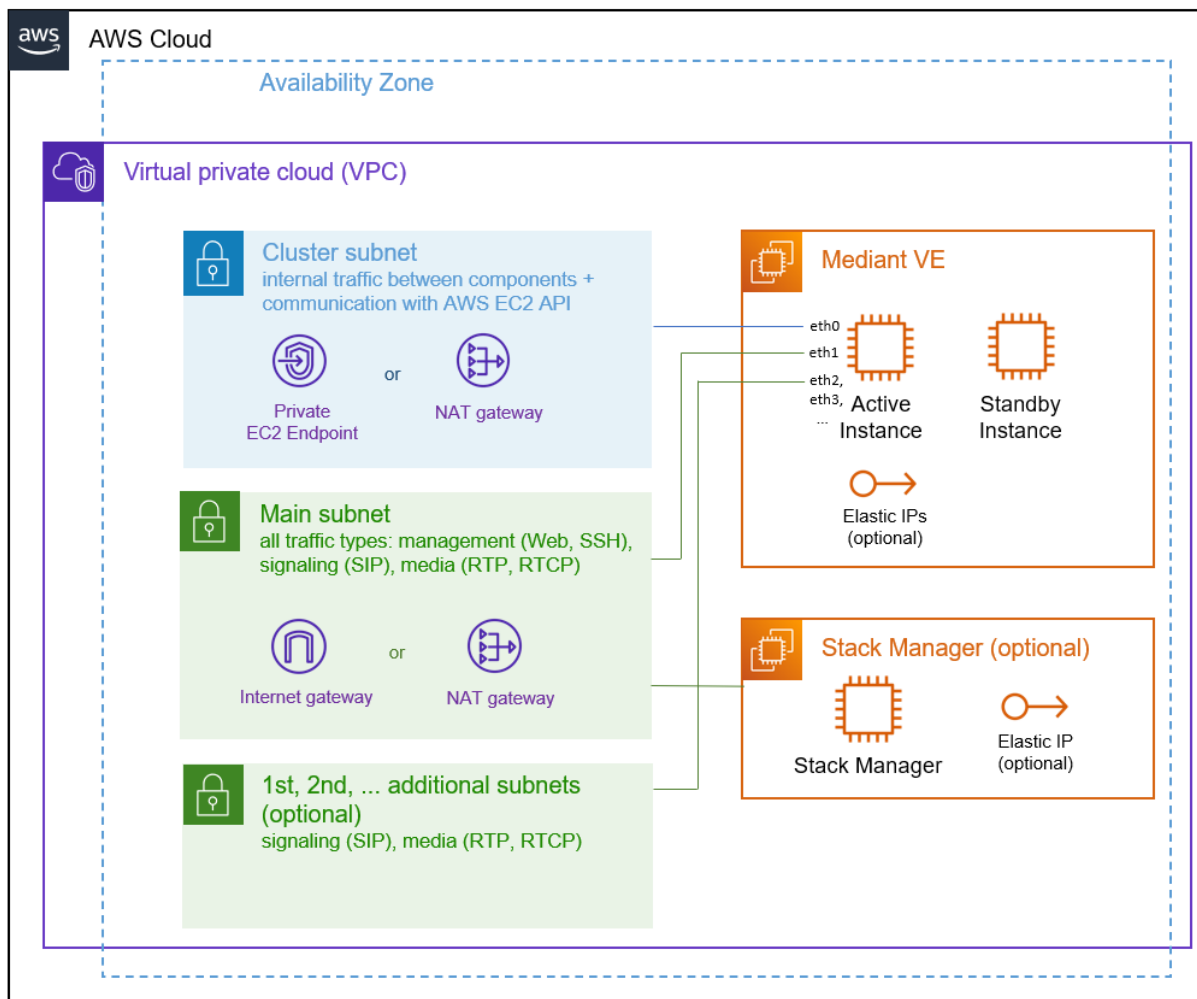


Mediant VE instance may be connected to multiple subnets, that must reside in the same Availability Zone of the Virtual Private Cloud (VPC) and be created prior to the Mediant VE deployment. Elastic IPs may be assigned to relevant network interfaces to enable communication via the public addresses.

2.2 HA Deployment in a Single Availability Zone

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

Figure 2-2: Network Architecture for HA Deployment in a Single Availability Zone



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

- **HA Subnet:** Carries internal communication between Mediant VE components. It is connected as the first network interface (eth0). For security reasons it is recommended to create dedicated HA subnet and protect it from access by other instances / equipment.
HA subnet is also used by active components 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 “HA Subnet” on page 29).
- **Main Subnet:** Carries management (HTTP, SSH, etc.), signaling (SIP) and media (RTP, RTCP) traffic; connected as the second network interface (eth1).
If deployment is performed via the **Stack Manager** it is also recommended to connect Stack Manager instance to the Main subnet, to enable seamless access from it to the Mediant VE’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.

- **1st, 2nd, ... Additional Subnets:** Carry signaling (SIP) and media (RTP, RTCP) traffic; connected as the third, fourth etc. network interfaces (eth2, eth3, ...) correspondingly; these subnets are optional, as the Main Subnet may carry all types of traffic.

Two instances are deployed and operate in 1+1 Active/Standby mode. Instances are connected to two (HA and Main) or more subnets. All subnets must reside in the same Availability Zone of the Virtual Private Cloud (VPC) and be created prior to the Mediant VE deployment.

Active instance uses secondary addresses on all network interfaces (except for the 1st one, connected to HA subnet) to communicate with external entities. During activity switchover these addresses are reassigned to another (newly active) instance by communicating with AWS EC2 API via the HA subnet.

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

If deployment is performed via the **Stack Manager**, 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 VE via private IP addresses. Elastic IP may be assigned to it to allow outbound access to AWS APIs and inbound access from the Internet.

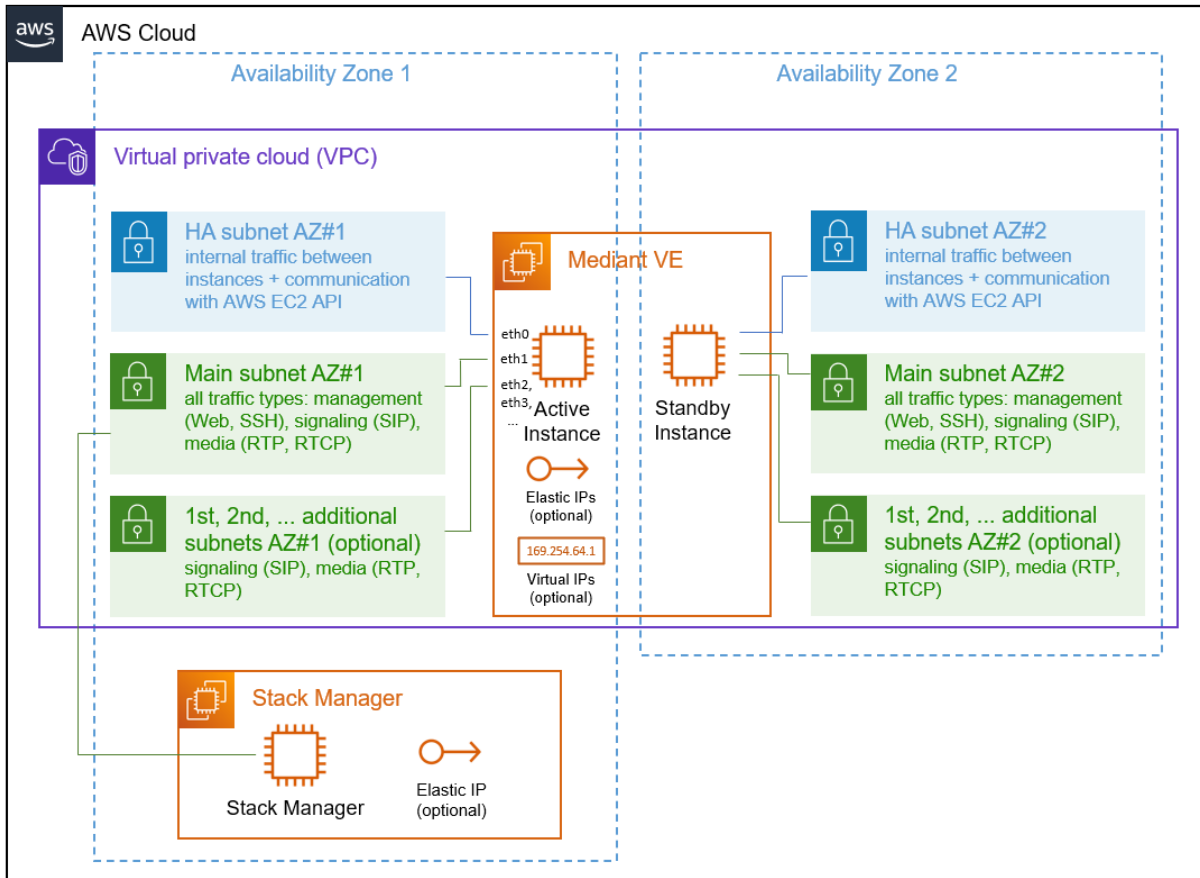
2.3 HA Deployment in Multiple Availability Zone



Note: HA deployment in multiple Availability Zones is supported starting from Version 7.4.500.

The following diagram shows the network architecture for High-Available Mediant VE deployment in multiple Availability Zones.

Figure 2-3: Network Architecture for HA Deployment in a Single Availability Zone



The deployment requires two sets of subnets – HA, Main and optionally Additional 1, Additional 2, etc. – that must be defined in two different Availability Zones. Two instances are deployed and operate in 1+1 Active/Standby mode. Each deployed instance resides in a different Availability Zone and is connected to the corresponding subnets set.

Each Mediant VE instance 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 the 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 2.3.1.

Reassignment of Elastic and Virtual IP addresses is performed by newly active component by communicating with AWS EC2 API via the HA subnet.



Note: Mediant VE HA deployment in multiple availability zones may use AWS Network Load Balancer instead or Virtual and/or Elastic IP addresses. For a detailed description, see Section 2.3.2'AWS Network Load Balancer.

2.3.1 Virtual IP Addresses

Mediant VE HA deployment in multiple availability zones uses Virtual IP addresses to enable communication between deployed Mediant VE SBC 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 the 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 via the `virtual_ip_ethX` advanced configuration parameter, for example:

```
virtual_ip_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 a 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 SBC instances. The entry is for a specific IP address (prefix /32) and is initially configured to point to the network interface of the 1st SBC instance that is initially active.

Figure 2-4: Virtual IP Address in AWS Route Table

The screenshot shows the AWS console interface for a route table (rtb-379b7d5e). The 'Routes' tab is selected, showing a list of 5 routes. The route for destination 169.254.64.1/32 is highlighted with a red box. The target for this route is the network interface eni-0c36cc91ef8ffeb81. Other routes include local destinations and VPC subnets (0.0.0.0/0, 20.0.0.0/24, 172.31.0.0/16).

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 SBC instance updates the entry with its own network interface, thus making sure that all communication via the Virtual IP address is sent towards it.

Figure 2-5: Virtual IP Address in AWS Route Table After the Switchover

The screenshot shows the AWS console interface for a route table (rtb-379b7d5e). The 'Routes (5)' section is active, displaying a list of routes. The route for destination 169.254.64.1/32 is highlighted with a red box, showing its target as eni-02cc5638a6de644ed. All routes are in an 'Active' state.

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 VE’s IP Interfaces table, and applications (e.g., SIP Interface) are “bound” to them.

Figure 2-6: Virtual IP Address in Mediant VE IP Interfaces Table

The screenshot shows the 'IP Interfaces (4)' table in the Mediant VE configuration. It lists four interfaces: eth0, eth1, eth2, and eth2:200. The eth2 interface is highlighted in blue and has the IP address 169.254.64.1.

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 SBC instances. 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.

2.3.1.1 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, as described in the previous chapter, ensures that other equipment deployed in the same AWS subnets (or in subnets with the same route tables) can communicate with the deployed Mediant VE SBC 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 VE SBC, it can use Virtual IP addresses to communicate with them.

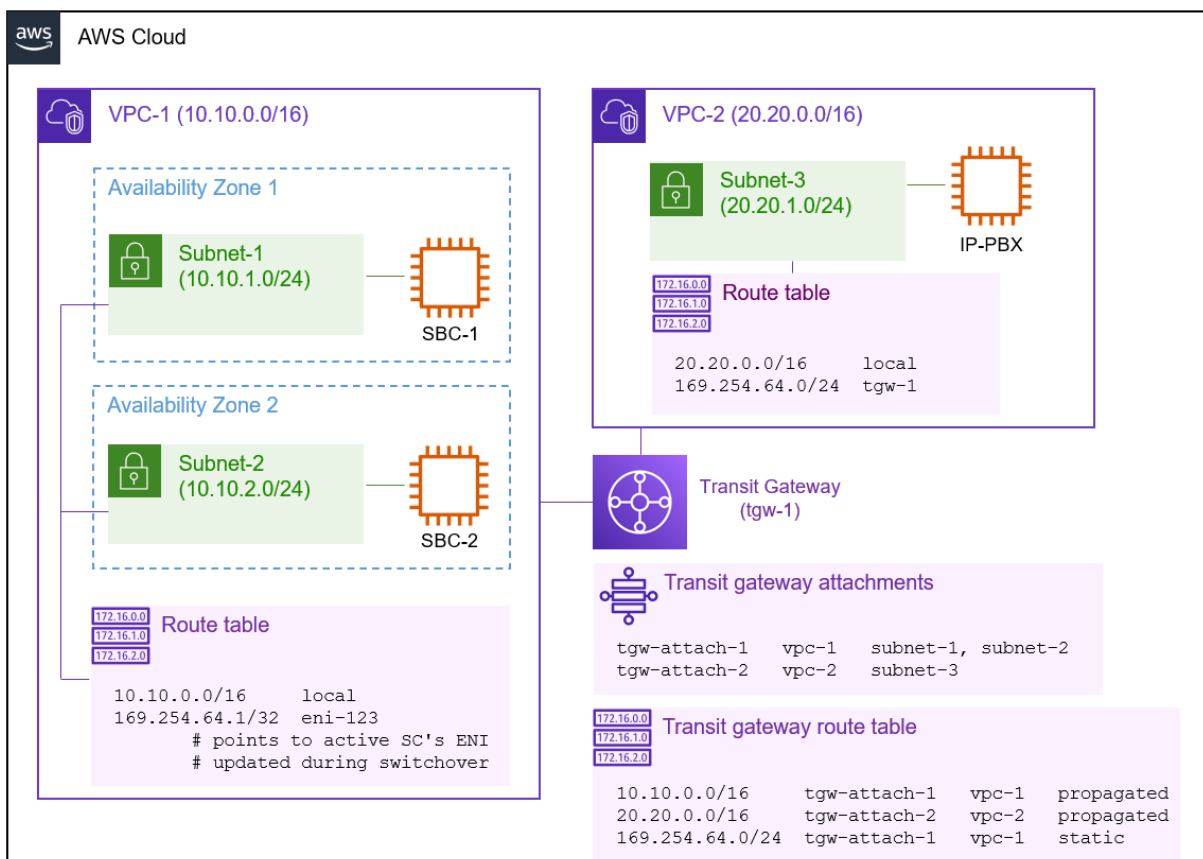
If equipment (e.g., IP-PBX or Contact Center) that needs to communicate with the Mediant VE SBC resides in a different AWS VPC or in the on-premise network, connected via AWS Direct Connect or site-to-site VPN, you must use **AWS Transit Gateway** to establish proper connectivity with Mediant VE SBC 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 2-7: Connection via Virtual IP Addresses across VPCs through 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 2-8: Virtual IP Addresses in AWS Transit Gateway Route Table

<input type="checkbox"/>	CIDR	Attachment ID	Resource ID	Resource type	Route type	Route state
<input type="checkbox"/>	100.0.0.0/24	tgw-attach-0311181a3b5380f83	vpc-05fe37c58e4430716	VPC	Propagated	Active
<input type="checkbox"/>	169.254.64.0/24	tgw-attach-0d981a54f2459e027	vpc-45f3152c	VPC	Static	Active
<input type="checkbox"/>	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 needs to communicate with Mediant VE resides, so that they route Virtual IP addresses (a specific one or the whole range) to 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 the 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. In the **Details** section, 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.

2.3.2 AWS Network Load Balancer

Mediant VE HA 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 VE HA deployment in multiple availability zones. It also mentions pros and cons of this deployment method.

The following operational modes are supported:

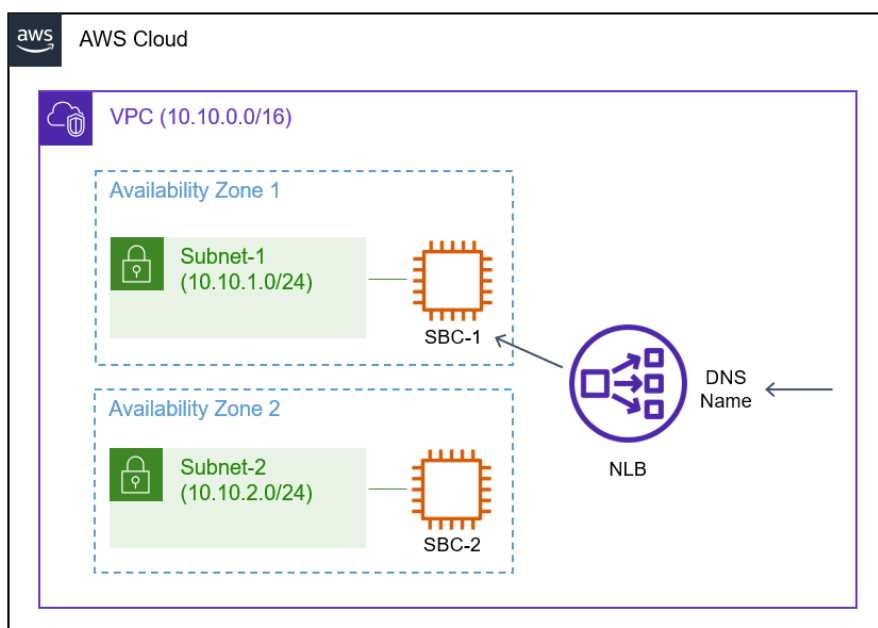
- NLB is not used. This is the default mode.
- Internal NLB is used instead of Virtual IP addresses. This is the alternative method, which is recommended.
- Public NLB is used instead of Elastic IP addresses.
- Internal / Public NLBs are used instead of both Virtual and Elastic IP addresses.
- Internal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

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

In a typical Mediant VE HA 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 status of the SBC instances that reside behind it via the keep-alive probes. At any time, one (active) SBC instance is "alive" and the other (standby) SBC instance is not. Therefore, NLB sends the received traffic to the active SBC instance. When an HA switchover happens, the newly active SBC instance becomes "alive" and NLB starts sending the traffic to it.

Figure 2-9: 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 front-end address. This DNS Name is resolved into one of the two different IP addresses (of internal NLB nodes). In other words, equipment that communicates with the Mediant VE 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 local IP addresses of the active SBC node and are reattached to the newly active SBC node (via SIP re-INVITEs) during the switchover.

Known limitation: 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 front-end DNS Name / FQDN is resolved into both A (IPv4) and AAAA (IPv6) records, while still using IPv4 addresses to communicate with the backend SBC 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 and 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 the 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 status of SBC instances and sends traffic to the active SBC instance accordingly.
 - Elastic IP addresses and Virtual IP addresses are reattached the active SBC instance via AWS APIs. These APIs may exhibit delays when the datacenter is overloaded.
- **Interoperability with SIP and 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 SIP / management equipment that communicates with the SBC the use of DNS Name and is able to refresh the address resolution in a timely manner. It also requires that external SIP equipment supports reattaching of media streams via SIP re-INVITEs, since the latter doesn't flow via the AWS NLB, but uses local IP addresses of the active SBC instance.
 - Elastic and Virtual IP addresses are moved during the switchover. Therefore, SIP / management equipment that communicates with the SBC works with a single set of IP addresses, regardless of which SBC 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 the 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.

2.3.2.1 Prerequisites for AWS Network Load Balancer Deployment

If you use AWS Network Load Balancer (NLB) for Mediant VE HA deployment, update the IAM role assigned to the 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 during creation of Mediant VE HA 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 via 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 the 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"}}
}
```

2.3.2.2 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_nlb = all - use internal / public NLBs instead of
Virtual / Elastic IP addresses
ha_nlb = oam - use internal NLB instead of Virtual IPs
for management traffic only
```

2.3.2.3 Configuration for AWS Network Load Balancer Deployment

When you use AWS Network Load Balancer (NLB) for Mediant VE HA 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 Network Load Balancer DNS Name / FQDN.

The latter can be found in the AWS dashboard or **More > Show IP Addresses** action in 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 through the NLB and therefore, always reach the active SBC instance, regardless of which SBC instance was active during the call establishment.



Note: Failure to configure the 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 an HA switchover.

2.3.2.4 Mediant VE Behavior During Switchover

A Mediant VE HA multi-zone deployment that uses AWS Network Load Balancer (NLB) is automatically configured to relatch media streams during a switchover. This is because media streams don't flow through the AWS NLB, but use the local IP addresses of the active SBC instance instead. Therefore, they need to be updated with new IP addresses upon the switchover.

This new switchover behavior applies to all SIP sessions, regardless of whether they traverse AWS NLB or not. It's not applicable to the "ha_nlb = oam" mode, where AWS NLB is used for management traffic only (e.g., to communicate with OVOC).

3 Prerequisites

Prior to deploying Mediant VE SBC on Amazon AWS, make sure that you meet the following prerequisites:

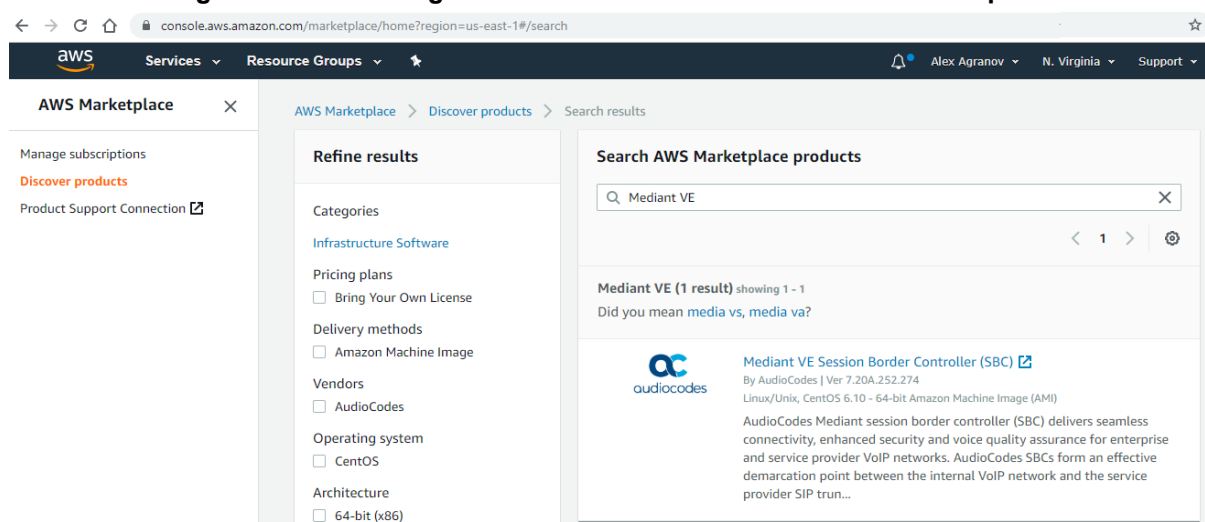
- You have an AWS account. If you don't have an AWS account, you can sign up for one on Amazon's website at <http://aws.amazon.com/>.
- You have subscribed to the AudioCodes Mediant VE offer in AWS Marketplace. See Section Subscribing to AudioCodes Mediant VE Product in AWS Marketplace for more information.
- You have created all subnets needed for Mediant VE deployment. See Section Network for more information.
- **For HA deployment:**
 - If you are going to perform deployment via CloudFormation template, make sure that you have received Mediant VE CloudFormation Template that is distributed as part of *Mediant VE Installation Kit*. See Section CloudFormation Template for Mediant VE HA Deployment for more information.
 - You have created an Identity and Access Management (IAM) role that enables Mediant VE to manage its network interfaces. See Section IAM Role for Mediant VE HA Deployment for more information.
 - You have created an HA subnet that is used for internal communication between Mediant VE instances and for accessing the AWS API during the activity switchover. See Section HA Subnet for more information.

3.1 Subscribing to AudioCodes Mediant VE Product in AWS Marketplace

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

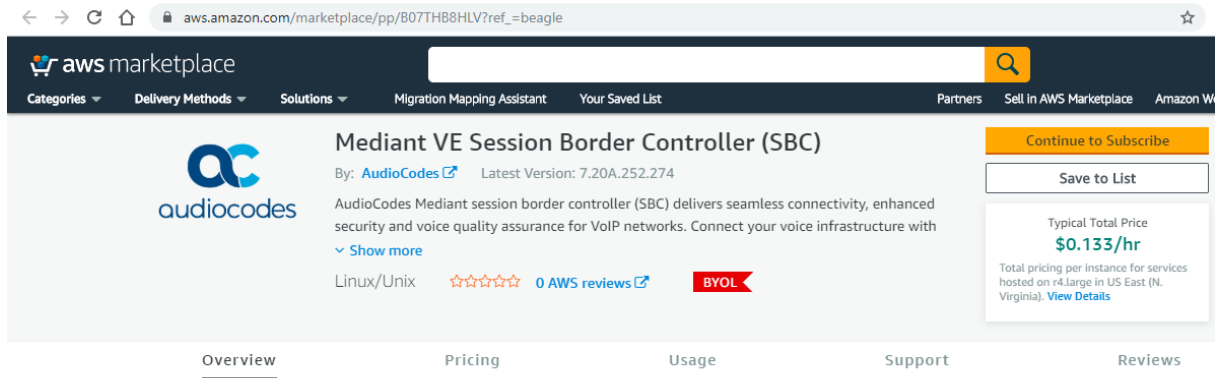
5. Open the AWS Marketplace console at <https://console.aws.amazon.com/marketplace>.
6. In the **Discover Products** tab, search for the "Mediant VE" product.

Figure 3-1: Searching for Mediant VE Product in the AWS Marketplace



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

Figure 3-2: 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

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

3.2 CloudFormation Template for Mediant VE HA Deployment

The CloudFormation template for single-zone High-Availability (HA) Mediant VE deployment is distributed as part of the *Mediant VE Installation Kit*, which can be downloaded from https://services.audiocodes.com/app/answers/detail/a_id/8.

Two CloudFormation templates are included:

- **sbc_ha_cloudformation.txt**: For regular AWS regions
- **sbc_ha_cloudformation_cn.txt**: For AWS China regions, for example, cn-north-1 and cn-northwest-1

3.3 IAM Role for Mediant VE HA Deployment

For HA deployment, the following IAM role must be created prior to deploying the Mediant VE instance. This role ensures that Mediant VE can manage its network interfaces and re-assign IP addresses during a switchover.

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



Note: IAM Role described below is needed only for HA deployment of Mediant VE, as described in the following sections:

Deploying High-Availability (HA) Mediant VE via CloudFormation Service and Deploying Mediant VE via Stack Manager. It is not needed for standalone deployment of Mediant VE, as described in Section Deploying Standalone Mediant VE via AWS EC2 Console.

➤ IAM Role for Single-Zone HA Mediant VE 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 HA Mediant VE 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:**

9. Open the AWS IAM console (<https://console.aws.amazon.com/iam>).
10. Navigate to the **Policies** screen, and then:
 - e. Click **Create**.
 - f. Select the **JSON** tab, copy-and-paste the IAM policy rules listed above, and then click **Review policy**.
 - g. Enter the IAM policy name (e.g., "SBC_HA"), and then click **Create policy**.
11. Navigate to the **Roles** screen, and then:
 - 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**.

3.3.1 IAM Role for Initial Configuration from S3 URL

Mediant VE SBC may be provided with an initial configuration INI file, stored on AWS Simple Storage Service (S3), during its launch. This is done by including the **#s3-url** element in the instance user-data, as described in [Automatic Provisioning of Mediant VE-CE SBC via Cloud-Init Configuration Note](#).

If you use this option, add the following rules to the IAM Role created previously, to enable Mediant VE 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/*"
    }
  ]
}

```

3.4 Network Architecture and Prerequisites

Mediant VE on Amazon Web Services (AWS) uses the following subnets:

- **Main Subnet:** Carries management (e.g. HTTP and SSH), signaling (SIP) and media (RTP, RTCP) traffic.
- **Additional Subnets:** Carry signaling (SIP) and media (RTP, RTCP) traffic. These subnets are optional and may be omitted if your network architecture doesn't require them.
- **HA Subnet:** Used for HA deployment only. Carries internal communication between Mediant VE instances. It's also used for accessing the AWS EC2 Endpoint during the switchover. See the next section for detailed instructions on how to create the HA Subnet.
- For Multiple Zones HA deployments, two sets of subnets must be created – in each Availability Zone.

3.4.1 HA Subnet

The HA subnet is used in high-availability (HA) Mediant VE deployments for the following tasks:

- Internal communication between Mediant VE instances
- Accessing AWS EC2 API (for IP address reassignment during activity switchover)

Mediant VE uses private addresses in the HA subnet. Therefore, to enable Mediant VE to access the AWS EC2 API via the HA subnet, you must do one of the following:

- (Recommended Method) Create a private EC2 endpoint in the HA subnet. This method creates a private AWS API endpoint inside the HA subnet, thereby enabling Mediant VE to access it via the private IP address.
- (Alternative Method) Attach a NAT gateway to the HA subnet. This method uses network address translation (performed by the NAT gateway) to enable access to public AWS EC2 API endpoint from Mediant VE SBC's private IP address.

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



Note: The following instructions describe how to create the HA subnet and enable access from it to the AWS EC2 API. Note that for Mediant VE deployment in multiple availability zones, you must create two HA subnets – one in each Availability Zone – and enable access to AWS EC2 API in each of them.

3.4.1.1 Creating HA Subnet

➤ **To create the HA 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. 'ha-route-table').
 - b. In the 'VPC' drop-down list, select the VPC where Mediant VE will be deployed.
 - c. Click **Create** to create the route table.

Figure 3-2: Creating Route Table

Route Tables > Create route table

Create route table

A route table specifies how packets are forwarded between the subnets within your VPC, the internet, and your VPN connection.

Name tag ⓘ

VPC* ↕ ⓘ

* Required

Cancel

3. Open the **Subnets** page, and then click **Create Subnet**.
 - a. In the 'Name tag' field, enter the new subnet name (e.g. 'ha-subnet').
 - b. From the 'Availability Zone' drop-down list, select the Availability Zone where Mediant VE 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 3-3: 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

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

Figure 3-4: Changing Cluster Subnet Route Table

Subnet: subnet-035888fc2f2e95bf8



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

Route Table: rtb-379b7d5e


Destination	Target
172.31.0.0/16	local
0.0.0.0/0	igw-0a49ae63

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

Figure 3-5: 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 **Save**



Note: Make sure that the HA subnet has a dedicated route table. Other subnets (i.e., Main subnet and Additional subnets) should be attached to different route table(s), that 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 .

3.4.1.2 Creating Private EC2 Endpoint in HA Subnet

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

➤ **To create the private EC2 endpoint in HA 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 VE 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 VE 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 select the 'Endpoint Security Group' that you created in the previous step.

- g. Click **Create Endpoint** to create the new endpoint.

Figure 3-6: 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 ⓘ

Filter by attributes 1 to 50 of more

Service Name	Owner	Type
com.amazonaws.eu-central-1.codebuild	amazon	Interface
com.amazonaws.eu-central-1.codecommit	amazon	Interface
com.amazonaws.eu-central-1.codepipeline	amazon	Interface
com.amazonaws.eu-central-1.config	amazon	Interface
com.amazonaws.eu-central-1.datasync	amazon	Interface
com.amazonaws.eu-central-1.dynamodb	amazon	Gateway
com.amazonaws.eu-central-1.ec2	amazon	Interface
com.amazonaws.eu-central-1.ecr.api	amazon	Interface
com.amazonaws.eu-central-1.transfer.server	amazon	Interface
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](#)

3.4.1.3 Creating NAT Gateway in HA Subnet



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

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

➤ **To create NAT Gateway and attach it to the HA 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 HA subnet was created (and where Mediant VE will be deployed) and that has an Internet Gateway attached to it. For example, select **Main Subnet**.



Note: Do not select **HA Subnet** at this stage. The NAT Gateway itself will be configured as a default route in the HA 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 3-7: 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 ↕ ↻ ℹ

Elastic IP Allocation ID* eipalloc-067ef98ad76079011 ↕ ↻

Create New EIP

New EIP (3.122.83.211) creation successful.



* Required

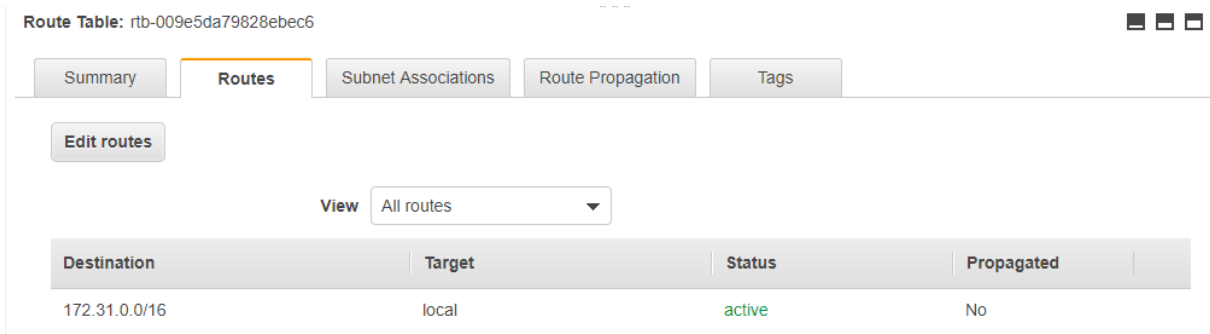
Cancel

Create a NAT Gateway

2. Open the **Route Tables** page, and then select the HA route table created in the previous steps.

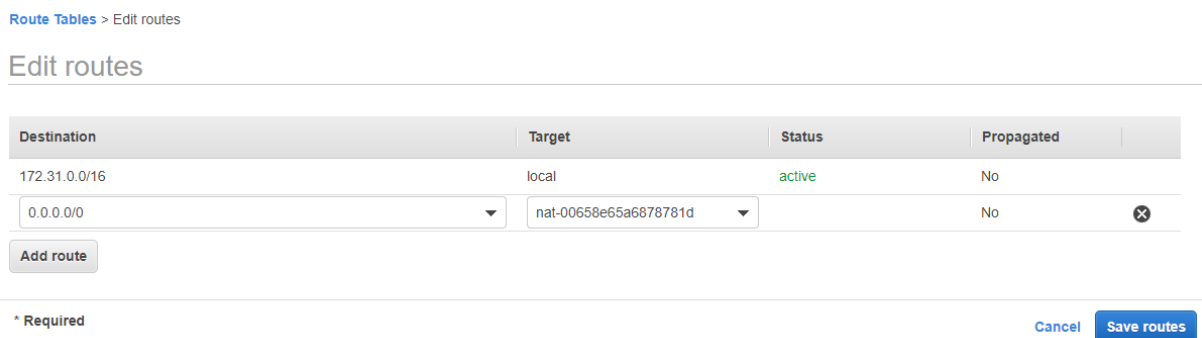
- Switch to the **Routes** tab, and then click **Edit routes** to edit the routes.

Figure 3-8: Editing Route Table



- 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 3-9: Creating Default Route



3.5 Instance Type

The following instance types are recommended for Mediant VE SBC deployment:

- For versions from 7.20CO, 7.40A and later streams based on OS Version 8:
 - **m5n.large**: This instance type is recommended for deployments that don't require transcoding and/or other DSP capabilities.
 - **c5n.2xlarge** or **c5n.9xlarge**: These instance types are recommended for deployments that require transcoding and/or other DSP capabilities.
- For versions from 7.20A stream based on OS Version 6:
 - **r4.large**: This instance type is recommended for deployments that don't require transcoding and/or other DSP capabilities.
 - **c4.2xlarge** or **c4.8xlarge**: These instance types are recommended for deployments that require transcoding and/or other DSP capabilities.

Refer to the [SBC Series Release Notes](#) for a complete list of instance types supported by Mediant VE SBC, their capacities and capabilities.

3.6 Automatic Configuration

Mediant VE SBC supports automatic configuration through the **cloud-init** mechanism. For more information, refer to the *Automatic Provisioning of Mediant VE SBC via Cloud-Init Configuration Note*.

This page is intentionally left blank.

4 Deploying Standalone Mediant VE via AWS EC2 Console

This section describes deployment for a standalone Mediant VE SBC via the AWS EC2 console.

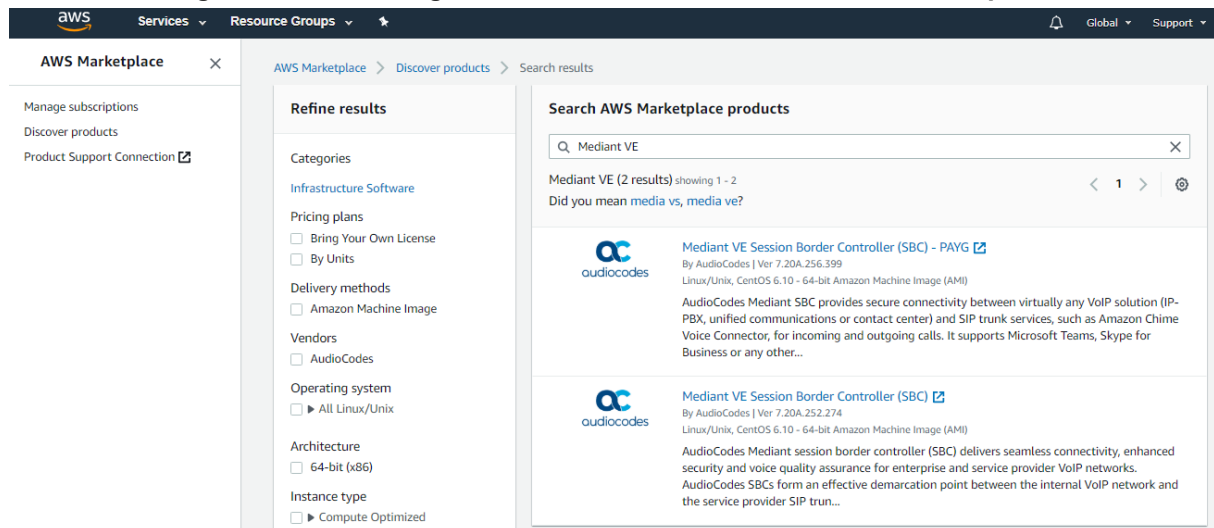


Note: This deployment method is applicable only to standalone (i.e., non-HA) deployments.

➤ **To deploy the standalone Mediant VE SBC instance:**

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 4-1: Searching for Mediant VE Product in the AWS Marketplace

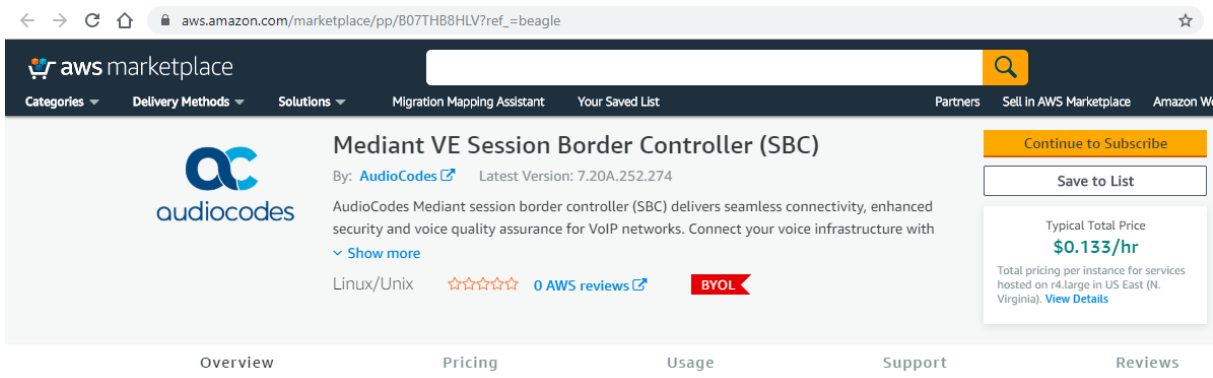


Two products are displayed:

- **Mediant VE Session Border Controller (SBC):** This product includes a trial license (limited to three SBC sessions) and requires a purchase of production license from AudioCodes.
- **Mediant VE Session Border Controller (SBC) – PAYG:** This product includes a pay-as-you-go license that enables Customers to use the SBC as much as needed and pay for the actual service consumed via their AWS account billing.

- Choose the Mediant VE product that matches your licensing needs. For example, choose **Mediant VE Session Border Controller (SBC)** product.

Figure 4-2: Mediant VE Product Page in AWS Marketplace



Product Overview

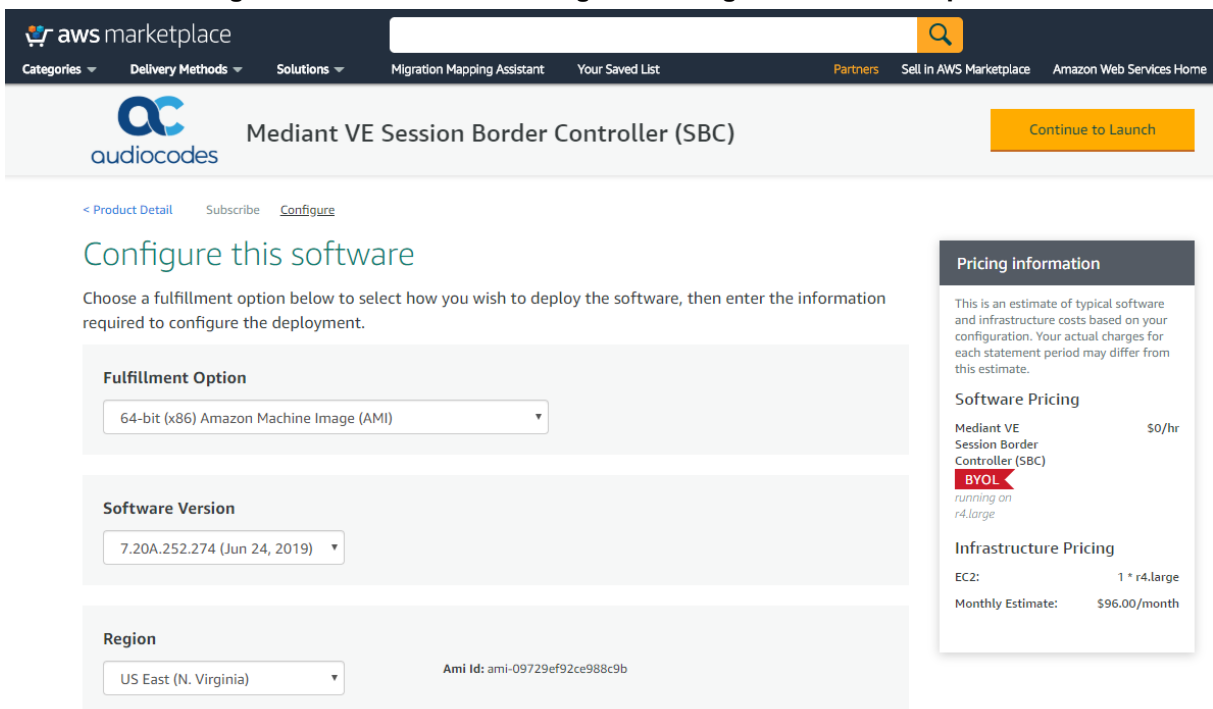
AudioCodes Mediant session border controller (SBC) delivers seamless connectivity, enhanced security and voice quality assurance for enterprise and service provider VoIP networks. AudioCodes SBCs form an effective demarcation point between the internal VoIP network and the service provider SIP trunk, 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

- Click **Continue to Subscribe** to subscribe to the Mediant VE SBC product.
- Click **Continue to Configuration** to proceed with SBC deployment.

Figure 4-3: Mediant VE Configuration Page in AWS Marketplace



6. Choose the software version that you want to deploy:
 - 7.20A stream is based on OS Version 6.
 - 7.20CO and 7.40A streams are based on OS Version 8 and provide significantly better performance and capacity (refer to the *SBC-Gateway Series Release Notes* for details).
7. Choose the Region where you want to launch the SBC.



Note: For the **Mediant VE SBC – PAYG** product, support is currently provided for installations in US regions only. For support in other regions, please contact AudioCodes at <https://online.audiocodes.com/aws-support-cloud>.

8. Click **Continue to Launch**.

Figure 4-4: Mediant VE Launch Page in AWS Marketplace

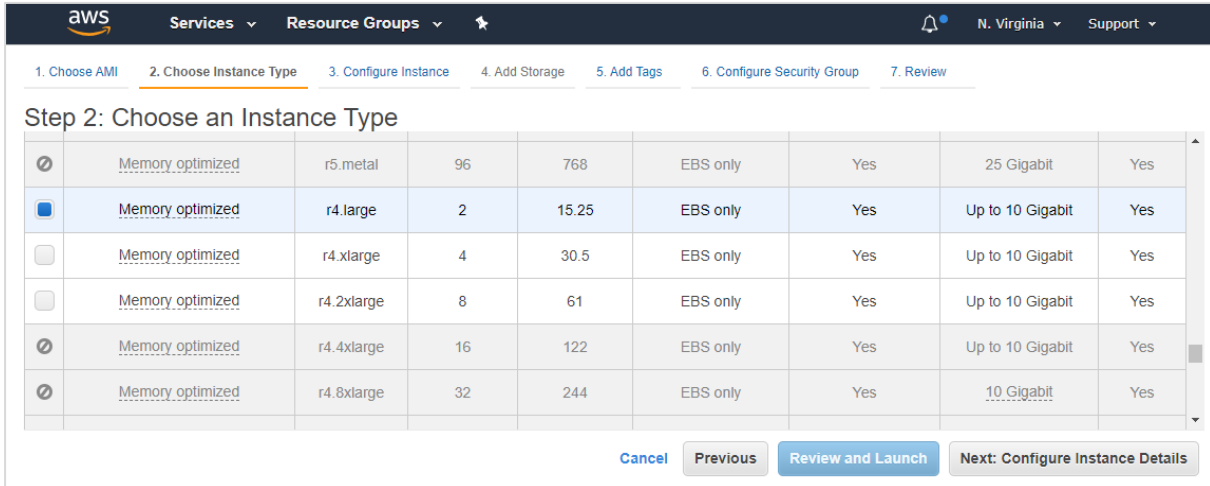
The screenshot displays the AWS Marketplace interface for the Mediant VE Session Border Controller (SBC). The header includes the AWS Marketplace logo and navigation links. The main content area features the AudioCodes logo and the product name. Below the product name, there are navigation links for Product Detail, Subscribe, Configure, and Launch. The main heading is "Launch this software". A sub-heading reads "Review your configuration and choose how you wish to launch the software." The "Configuration Details" section lists the following information:

Fulfillment Option	64-bit (x86) Amazon Machine Image (AMI) Mediant VE Session Border Controller (SBC) <i>running on r4.large</i>
Software Version	7.20A.252.274
Region	US East (N. Virginia)

Below the configuration details is a "Usage Instructions" button. The "Choose Action" section contains a dropdown menu with "Launch through EC2" selected. A note next to the dropdown states: "Choose this action to launch your configuration through the Amazon EC2 console." A prominent orange "Launch" button is located at the bottom right of the configuration area.

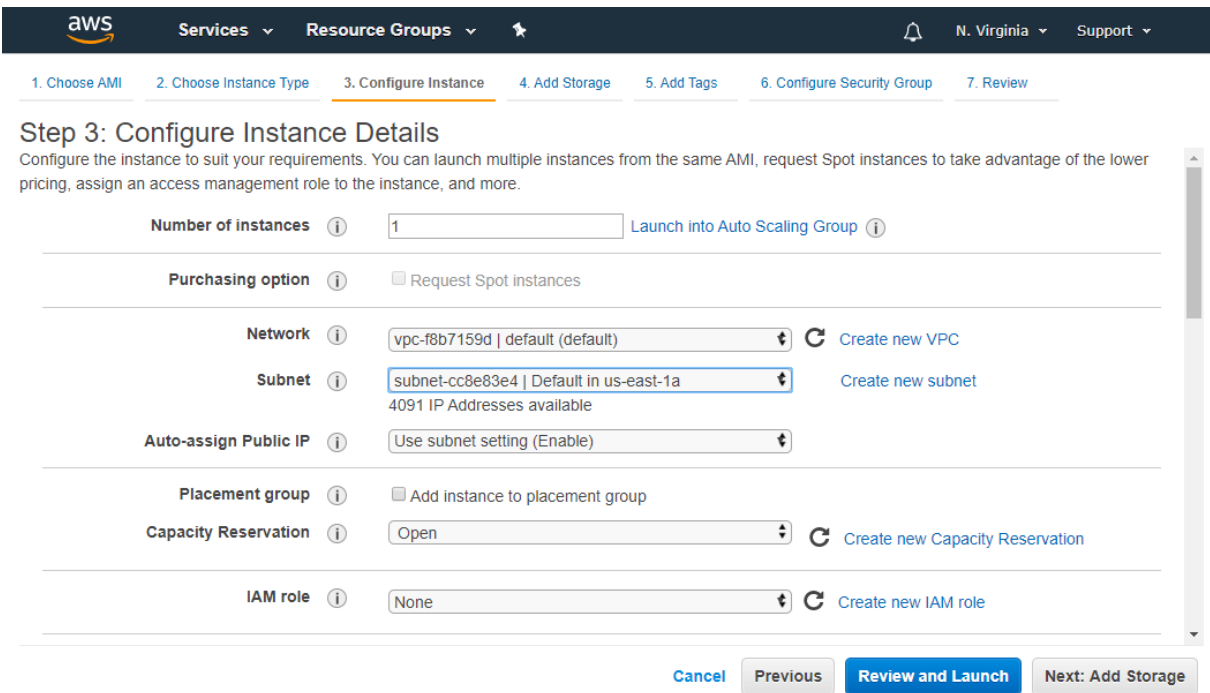
9. From the 'Choose Action' drop-down list, select **Launch through EC2**, and then click **Launch**; the Choose Instance Type page appears:

Figure 4-5: Choose Instance Type Page



10. Choose the instance type as described in Section Instance Type.
11. Click **Next**; the Configure Instance page appears:

Figure 4-6: Configure Instance Page



12. Configure network devices and IP addresses:
 - For **Network**, select the VPC where SBC should be deployed.
 - For **Subnet**, select the LAN Subnet. This subnet is used to communicate with the Enterprise IP-PBX and for accessing the SBC management interface (Web or CLI).
 - For **IAM role**:
 - ◆ If you are deploying the **Mediant VE SBC – PAYG** product, select Automatically create an IAM role with the required permission and the name below, and then enter the IAM role name (e.g., "metering-role").
 - ◆ If you are deploying the **Mediant VE SBC** product, leave IAM role empty.



Note: The **Mediant VE SBC – PAYG** product requires an IAM role with the following policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        aws-marketplace:MeterUsage
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

This role allows Mediant VE SBC PAYG instance to communicate with the AWS Metering API and must be assigned to the launched instance – either automatically (as described above) or manually.

- If you want the deployed instance to have multiple network interfaces, in the **Network Interfaces** section located at the bottom of the page, click **Add Device**, and then select the subnet for the added device (**eth1**).
- If you want the deployed instance to have multiple IP addresses on the same network interface, in the **Network Interfaces** section located at the bottom of the page, click **Add IP**.

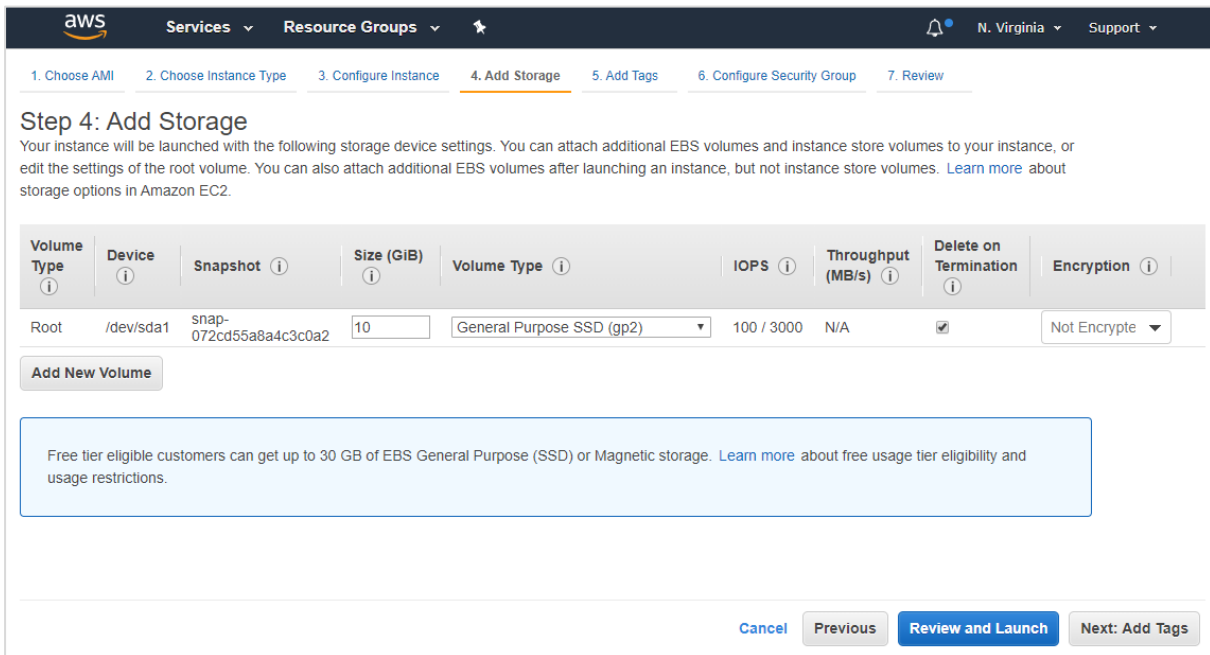


Notes:

- If your instance has only one network interface, AWS EC2 may automatically assign a public IP address to the instance. The exact behavior depends on the VPC and/or Subnet configuration. This address however changes if you stop/start the instance and therefore is typically not useful for production environment.
- If you configure multiple network interfaces, AWS EC2 does not automatically assign public IP addresses for the instance.
- To make the Mediant VE SBC instance properly reachable from the Internet, you should assign Elastic IP addresses to it, as described in Section Assigning Elastic IP Addresses to the Instance.
- AWS EC2 Web console supports the configuration of up to two network devices during instance launch. To overcome this limitation and define additional network devices during initial instance launch, consider using AWS EC2 CLI or AWS CloudFormation instead.
- If you need to add/remove network interfaces and/or secondary IP addresses after the initial instance deployment, do this through AWS EC2 management interfaces (i.e., AWS EC2 Web console or AWS EC2 CLI). Mediant VE SBC software automatically detects these changes and updates its networking configuration accordingly.

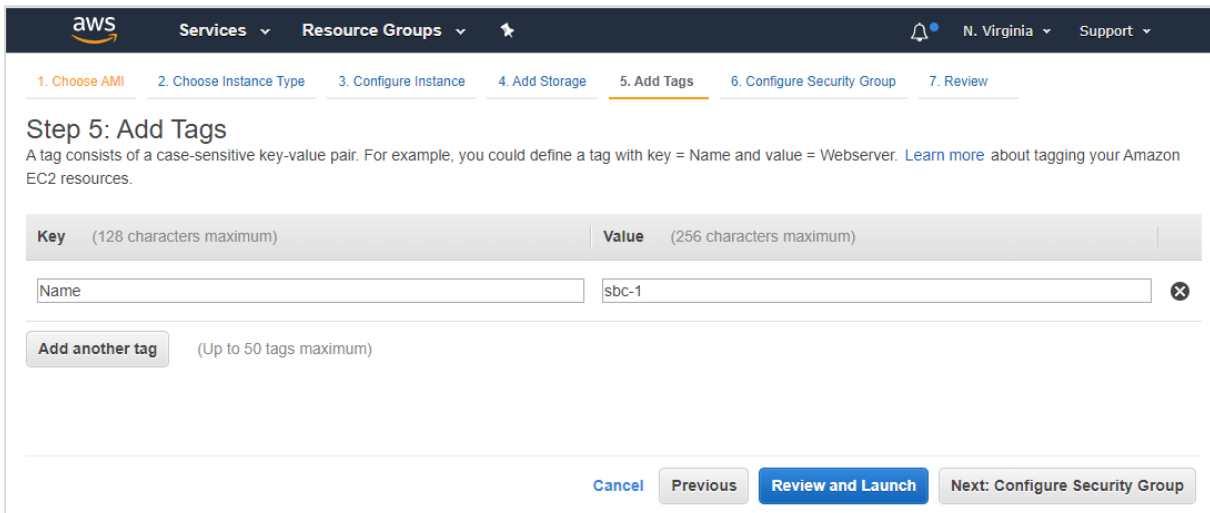
- Click **Next**; the Add Storage page appears:

Figure 4-7: Add Storage Page



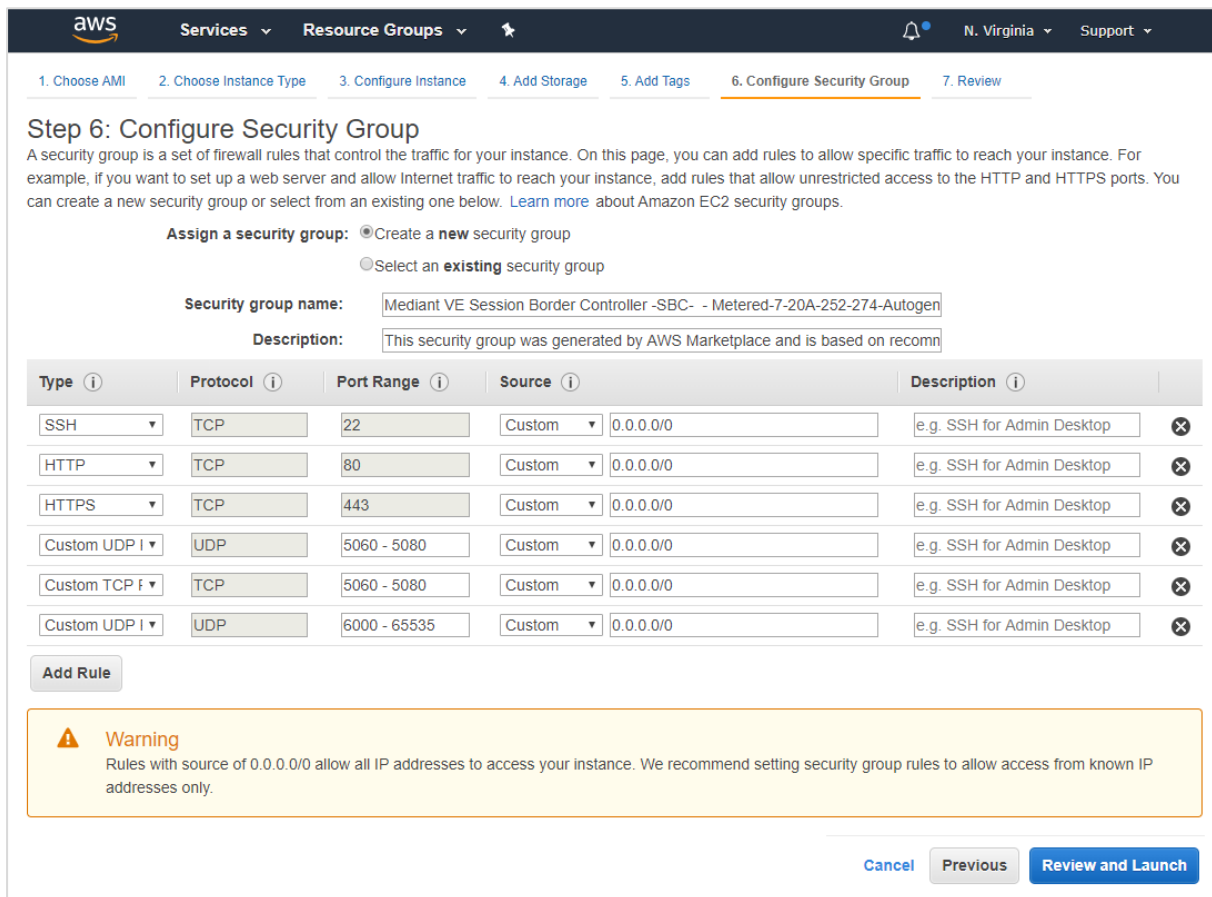
- From the 'Volume Type' drop-down list, select the required volume of the instance. This setting does not affect SBC performance and may be set to any value.
- Click **Next**; the Tag Instance page appears:

Figure 4-8: Tag Instance Page



- In the 'Value' field, enter a name for your instance, and then click **Next**; the Configure Security Group page appears:

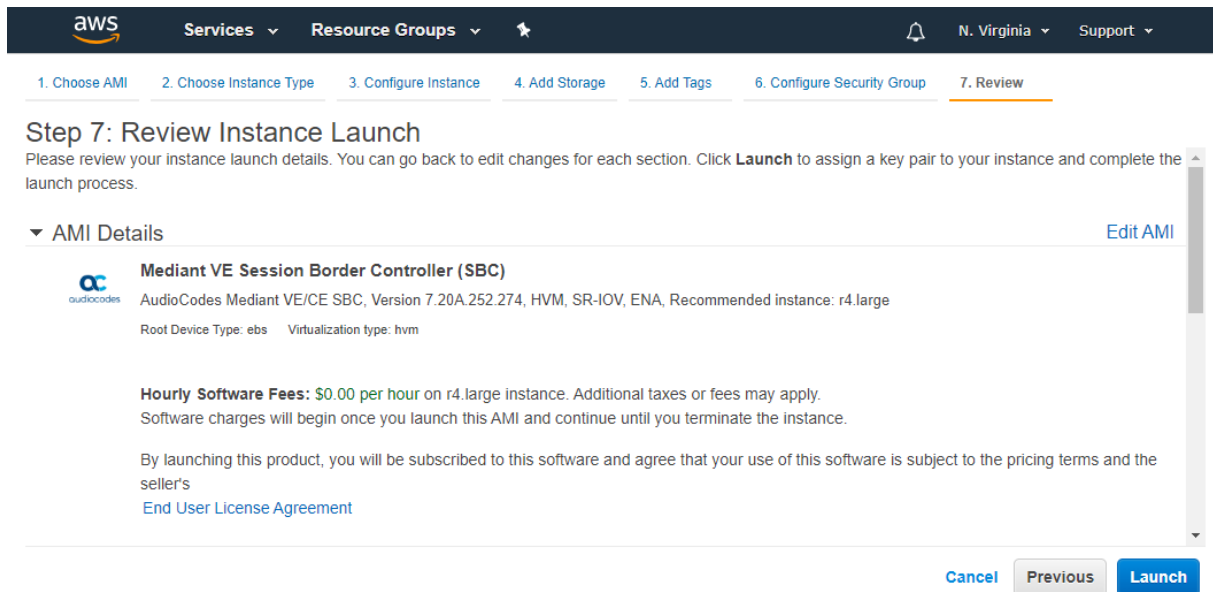
Figure 4-9: Configure Security Group Page



- Configure firewall rules to allow management (SSH, HTTP, and HTTPS), signaling (SIP) and media (RTP/RTCP) traffic with your instance. Use default rules as a starting point and modify them to match your actual deployment needs.

18. Click **Review and Launch**; the Review page appears displaying a summary of your instance configuration:

Figure 4-10: Review Page



19. Click **Launch**; the Select an existing key pair window appears.
20. Select a key pair to authenticate SSH connection with the SBC instance, click the **I acknowledge** check box, and then click **Launch Instances**.
21. Wait until the new Mediant VE instance is deployed and fully starts (it may take up to 5 minutes). Navigate to the **Instances** page and check the *instance-id* of the deployed instance.
22. Proceed to the next step to assign Elastic IPs to the launched SBC instance.
23. Once you're finished with networking configuration, log in to the deployed instance using the following default credentials:
 - Username: **Admin**
 - Password: *instance-id*

4.1 Assigning Elastic IP Addresses to the Instance

The AWS EC2 environment assigns “private” IP addresses to the instances running in it. These addresses may be used for communication between the instances running inside the same network (VPC); however, they may not be used to connect to the instance over the Internet.

If the instance has only one network device, AWS EC2 may automatically assign a public IP address to it. The exact behavior depends on the VPC and/or Subnet configuration. This address however is taken from a “shared pool” and changes if you stop/start the instance. Therefore, it is not very useful for production environment.

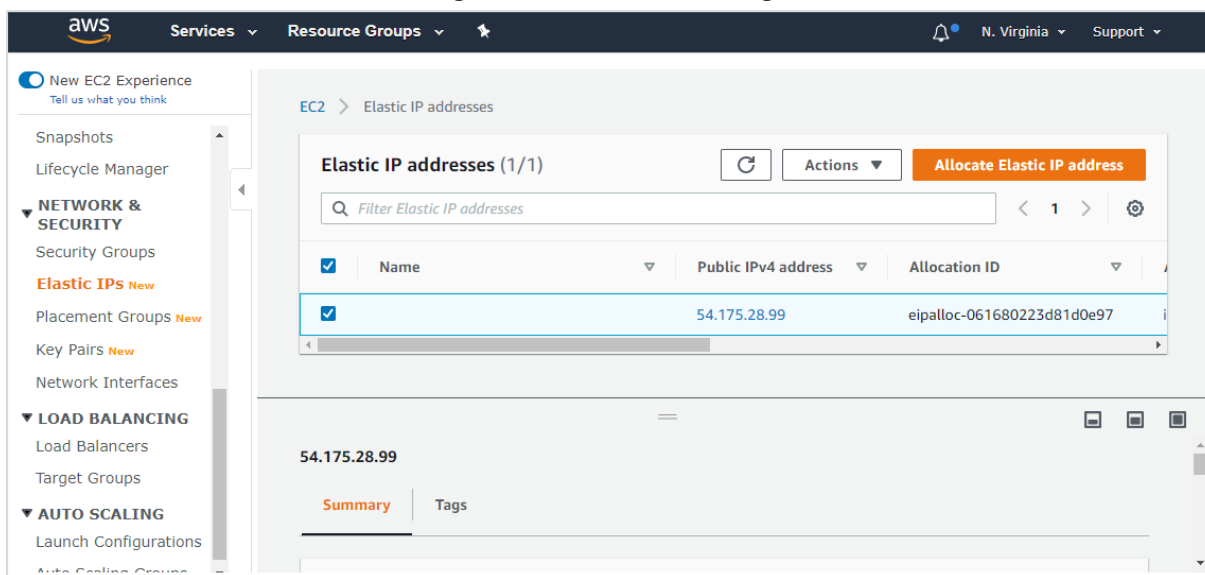
To make SBC properly reachable over the internet, you must allocate Elastic IP addresses and assign them to your instance. Multiple Elastic IP addresses may be assigned to the same AWS EC2 instance, depending on the number of configured private IP addresses.

When an Elastic IP address is associated with the specific instance’s private IP address, AWS EC2 environment performs NAT translation by converting elastic IP address to the private IP address, while preserving the port range. If the SBC needs to communicate with a SIP entity using the Elastic IP address, the latter must be configured in the NAT Translation table to ensure proper modification of SIP / SDP messages for NAT traversal.

➤ **To allocate Elastic IP address to SBC instance:**

1. Open the EC2 console at <https://console.aws.amazon.com/ec2>.
2. Navigate to the **Elastic IPs** page under NETWORK & SECURITY:

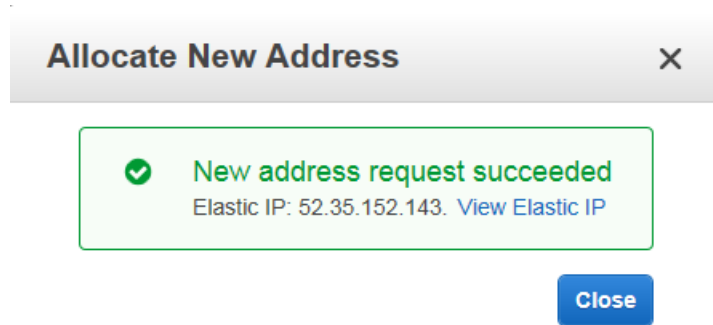
Figure 4-11: Elastic IPs Page



3. Click **Allocate New Address**; a message box appears requesting you to confirm.

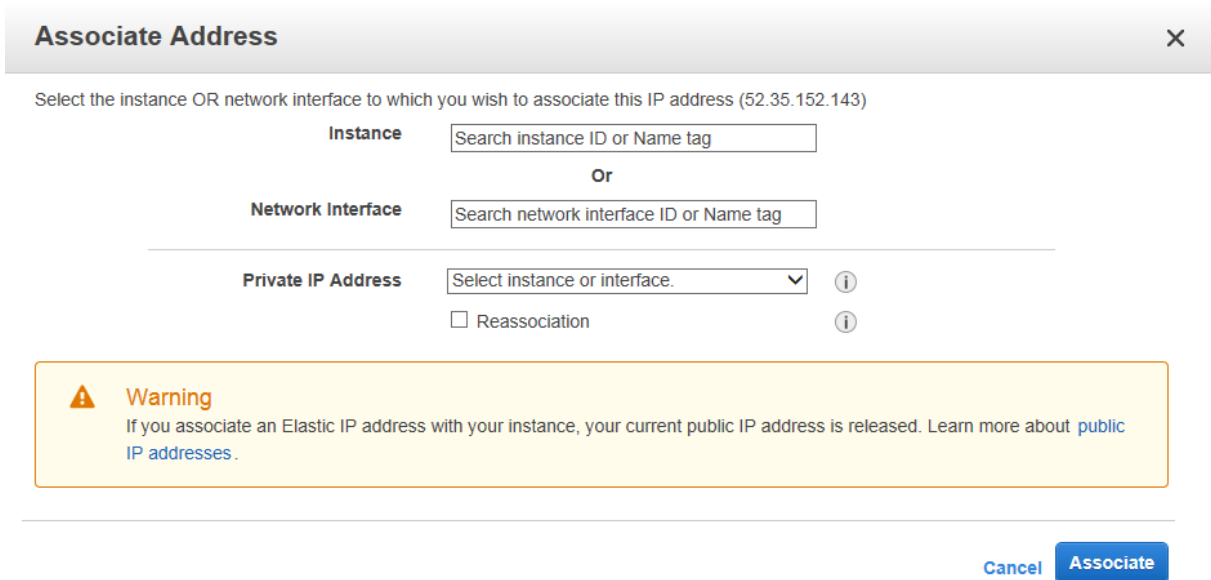
- Click **Yes, Allocate** to confirm; a message box appears displaying the allocated IP address:

Figure 4-12: Allocated IP Address



- Click **Close** to close the message box.
- From the Actions drop-down list, select **Associate Address**.

Figure 4-13: Associate Address Window



- Select the instance or network interface and private IP address to which you want to associate the Elastic IP address, and then click **Associate**.
- If you have configured multiple IP addresses and want to make them reachable over the Internet as well, repeat the procedure for additional IP addresses.

4.2 Post-Installation Configuration of Mediant VE SBC – PAYG Product

The "Mediant VE SBC – PAYG" product includes a pay-as-you-go (PAYG) license that requires a persistent connection between the SBC and the AWS Metering API. This connection is performed using public IP addresses. Therefore, you must assign an Elastic IP address to one of the Mediant VE SBC's network interfaces, as described in the previous section. You must also configure the Mediant VE SBC to use this interface, as described below.



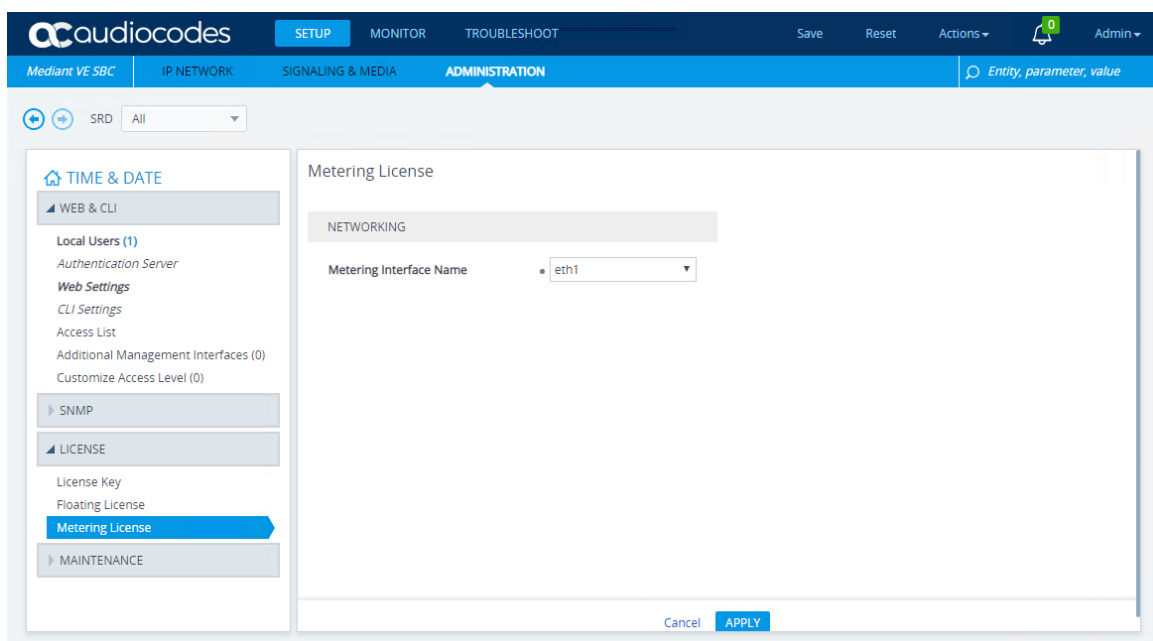
Notes:

- This section is applicable only to the "Mediant VE SBC – PAYG" product that uses the pay-as-you-go (PAYG) license.
- The following procedure is extremely important, as without it the SBC will be unable to communicate with the AWS Metering API and will **seize** its service.

➤ To perform post-installation configuration of Mediant VE SBC – PAYG product:

1. Open the SBC Web interface.
2. Log in using default credentials:
 - Username: **Admin**
 - Password: **instance-id**
3. Navigate to the Metering License page (**ADMINISTRATION > LICENSE > Metering License**).

Figure 4-14: Metering License Page



4. From the 'Metering Interface Name' drop-down list, select the network interface that has an Elastic IP address attached. This network interface will be used to communicate with the AWS Metering API.
5. Click **Apply** to apply your settings.
6. On the toolbar, click **Save** to save your settings.

4.2.1 Troubleshooting Mediant VE SBC – PAYG Deployment

Mediant VE SBC – PAYG deployment requires persistent connection between the SBC and AWS Metering API. If this connection is unavailable, SBC seizes its operation and raises the "No connection to Metering API" alarm.

SBC contacts the AWS Metering API after the call is completed. Therefore, you will not be able to detect if there is a problem until you have perform a few calls.

Typical reasons that may prevent proper connection between the SBC and the AWS Metering API include the following:

- No "metering" IAM role (with `aws-marketplace:MeterUsage` action) assigned to the SBC instance.
- No Elastic IP address on the network interface used for communication with the AWS Metering API (**ADMINISTRATION > LICENSE > Metering License**).
- A Network Security Group or some other firewall device is blocking communication between the SBC and the AWS Metering API endpoint (`https://metering.marketplace.<region>.amazonaws.com`).
- Incorrect UTC time configured on the SBC (**ADMINISTRATION > TIME & DATE**).

While troubleshooting, it may be useful to enable the detailed metering logs, by connecting to the SBC's CLI interface and issuing the following commands:

```
enable
  <password> (default: Admin)
debug license aws-payg toggle-cl-debug
```

Collect the logged messages via one of the following means:

- Web interface's **TROUBLESHOOT > MESSAGE LOG** page
- CLI interface's `debug log` command
- Syslog Viewer utility available at <https://tools.audiocodes.com/install/>

You may need to wait up to 15 minutes for the metering logs to accumulate after enabling them.

5 Deploying High-Availability (HA) Mediant VE via CloudFormation Service

This section describes deployment of high-availability (HA) Mediant VE that includes two EC2 instances, operating in 1+1 Active/Standby mode. Both instances are connected to the same network subnets and deployed into a single Availability Zone. See Section HA Deployment in a Single Availability Zone for deployment topology description.

**Note:**

- This deployment method is applicable only to "single zone" HA deployment topology.
- HA deployment is supported only by the **Mediant VE SBC** product (and not by the **Mediant VE SBC – PAYG** product).

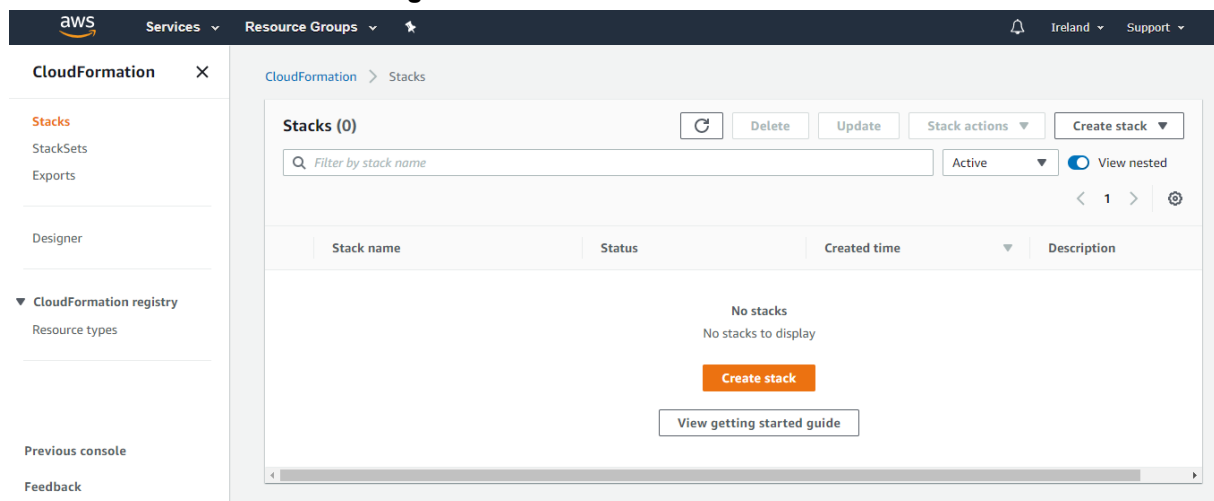
The deployment is performed via the CloudFormation service. The reference CloudFormation template can be downloaded from AudioCodes Services Portal at https://services.audiocodes.com/app/answers/detail/a_id/8 as part of the *Mediant VE Installation Kit*.

The CloudFormation template provided by AudioCodes has certain limitations. For example, it attaches the Elastic IP to the management interface of the deployed Mediant VE instance, but not to the additional interfaces (if used). Customers should use the provided CloudFormation as a reference and modify it to match their deployment needs.

➤ **To deploy high-availability (HA) Mediant VE via AWS CloudFormation service:**

1. Open the CloudFormation console at <https://console.aws.amazon.com/cloudformation>

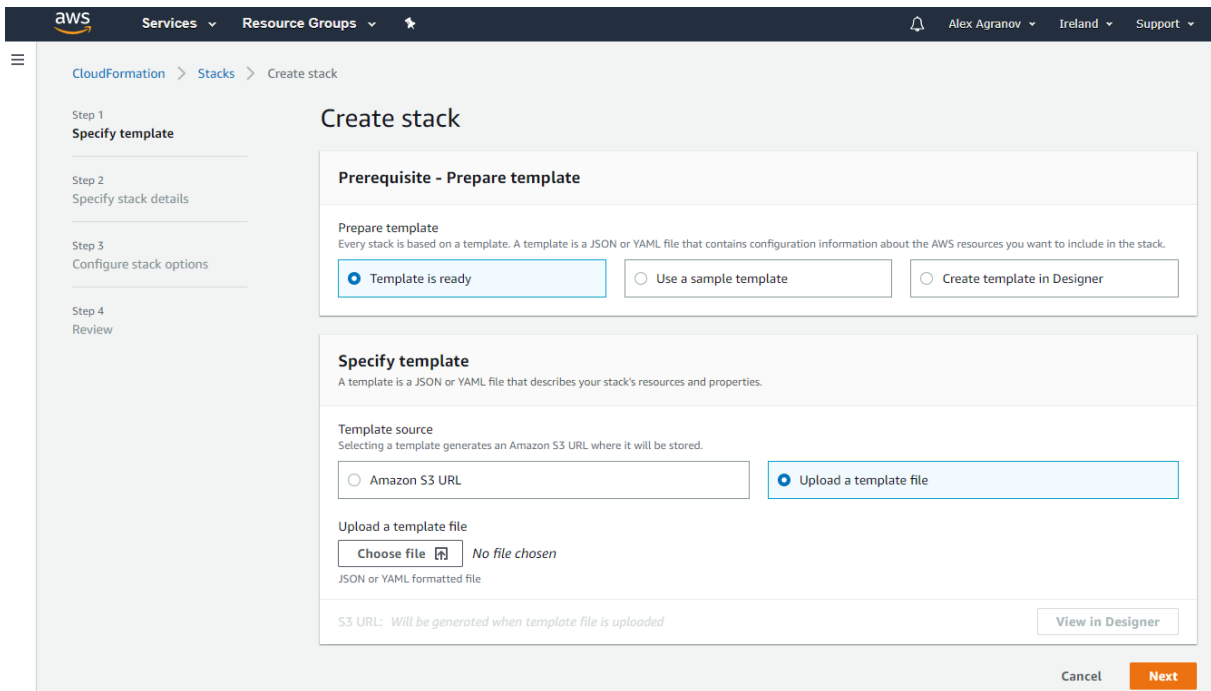
Figure 5-1: CloudFormation Console



2. Select the Region (in the upper right corner) in which to perform the deployment.

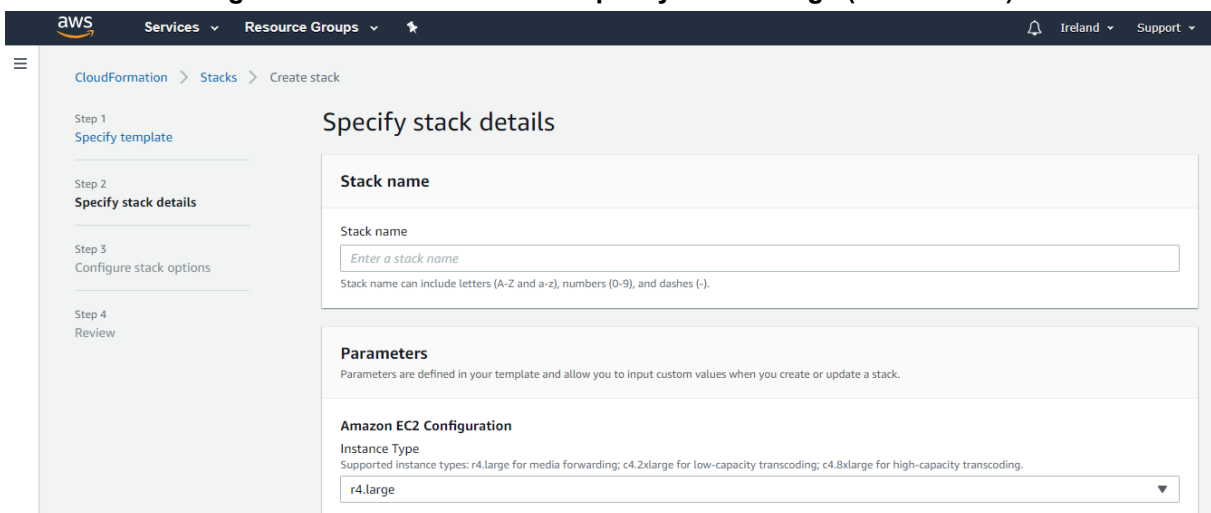
3. Click **Create Stack** to create a new stack, and then select **With new resources (standard)** from the drop-down menu; the Create Stack page appears:

Figure 5-2: CloudFormation – Create Stack Page



4. Under the **Specify template** group, select the **Upload a template file**, click **Choose File**, and then select the *Mediant VE HA CloudFormation template* file provided by AudioCodes.
5. Click **Next**; the Specify Stack Details page appears with the fields populated with parameter settings from the template file that you loaded in the previous step:

Figure 5-3: CloudFormation - Specify Details Page (Stack Name)



6. In the **Stack Name** field, type in a meaningful stack name. The stack name is an identifier that helps you find a particular stack from a list of stacks. A stack name can contain only alphanumeric characters (case-sensitive) and hyphens. It must start with an alphabetic character and can't be longer than 128 characters.

7. Under the **Parameters** section, configure parameters to match the desired stack configuration:
 - Amazon EC2 Configuration:
 - ◆ **Instance type:** AWS EC2 instance type for the stack.
 - ◆ **Amazon Machine Image (AMI):** Amazon Machine Image (AMI) ID of Mediant VE SBC (check the Mediant VE product in the AWS Marketplace to find AMI ID for the specific region).
 - ◆ **IAM Role:** Name of the existing IAM role that enables Mediant VE to manage its network interface, as created in Section IAM Role for Mediant VE HA Deployment.
 - ◆ **Key Name:** Name of the existing Key Pair used to secure access to the Mediant VE's SSH interface.
 - ◆ **S3 URL of INI Configuration File:** (Optional) Amazon S3 URL of initial Mediant VE configuration file.




Note: If you configure a value for “S3 URL of INI Configuration File”, make sure that the IAM role allows access to the corresponding S3 bucket, as described in Section IAM Role for Initial Configuration from S3 URL.

- Network Configuration:
 - ◆ **Which VPC should the SBC be deployed to?** VPC ID of the existing Amazon Virtual Private Cloud (VPC) where Mediant VE should be deployed.
 - ◆ **Number or Network Interfaces:** Number of network interfaces to be attached to Mediant VE SBC instances. Minimum number is 2; maximum number depends on the instance type used. Refer to Section Network for details.
 - ◆ **HA Subnet:** Subnet ID of existing subnet in your VPC. The subnet is used for internal traffic between two SBC instances and for accessing AWS API. The subnet must have a private EC2 API endpoint or a NAT Gateway set as default route, as described in Section HA Subnet. It is attached to the 1st network interface (eth0).
 - ◆ **Main Subnet:** Subnet ID of existing subnet in your VPC. The subnet is used for Management traffic (e.g., for accessing the SBC's Web interface). It may also be used for VoIP traffic (signaling and media). It is attached to the 2nd network interface (eth1).
 - ◆ **1st Additional Subnet:** Subnet ID of existing subnet in your VPC. The subnet is used for VoIP traffic (signaling and media). It is attached as the 3rd network interface (eth2). If 'Number of Network Interfaces' is less than 3, set this parameter to the same value as 'Main Subnet'.
 - ◆ **2nd Additional Subnet:** Subnet ID of existing subnet in your VPC. The subnet is used for VoIP traffic (signaling and media). It is attached as the 4th network interface (eth3). If 'Number of Network Interfaces' is less than 4, set this parameter to the same value as 'Main Subnet'.
 - ◆ **Public IPs:** Select which network interfaces should be assigned with public (Elastic) IP addresses. Keep in mind that for Elastic IPs to operate correctly, corresponding subnets must have an Internet Gateway set as the default route.
8. Click **Next**; the Options page appears. Leave this page at its default settings.

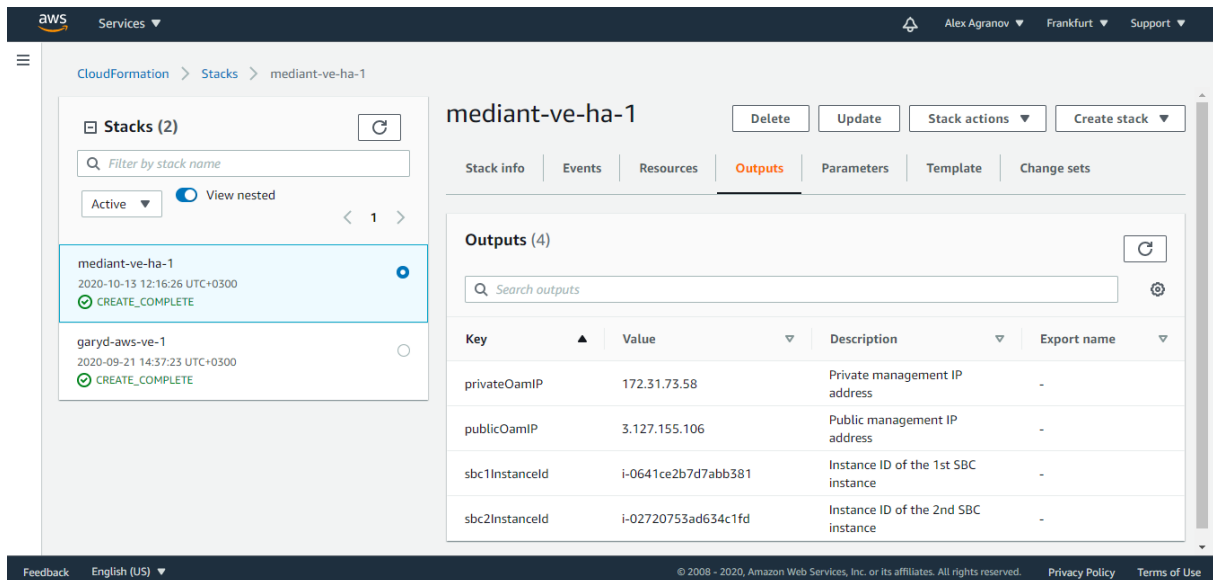
9. Click **Next**; the Review page appears, showing a summary of your stack settings:
10. Click **Create**; CloudFormation starts creating the stack. During stack creation, its state changes to "CREATE_IN_PROGRESS".

Figure 5-4: CloudFormation – Stack Creation Progress



11. Wait until the stack is created and its state changes to "CREATE_COMPLETE". Two SBC instances are created and configured to operate in 1+1 active/standby mode. Their instance-ids, management IPs and default admin credentials are listed in the **Outputs** tab.

Figure 5-5: CloudFormation – Stack Outputs



12. Access the SSH or Web interface of the deployed Mediant VE SBC using the IP address from the **privateOamIP** or **publicOamIP** field, listed in the **Outputs** tab. Use the default admin credentials – from **adminUsername** and **adminPassword** fields – to log in.



Note: If you copy/paste the *instance-id* from the **Outputs** tab, the browser may append a space to the copied value, thus making it invalid. Therefore, it is recommended to type *instance-id* manually.

5.1 Deleting HA Mediant VE Deployment

To delete deployed Mediant VE stack, use **Delete** action from the CloudFormation screen.

6 Deploying Mediant VE via Stack Manager

This section describes the deployment of Mediant VE via the Stack Manager.



Note: This method is applicable only to the **Mediant VE SBC** product (and not the **Mediant VE SBC – PAYG** product). Both standalone and HA deployment topologies are supported.

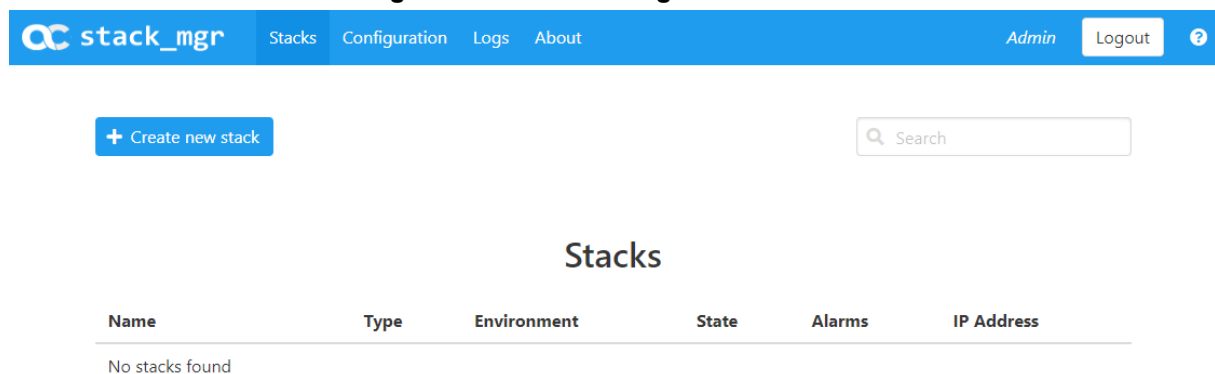
Stack Manager is a management tool developed by AudioCodes that enables simple and intuitive deployment and complete lifecycle management of Mediant VE and Mediant CE products on public clouds. The tool provides support for all Mediant VE deployment topologies – standalone, “single zone” HA and “multiple zones” HA – and may be used not only for initial product deployment, but also for “day 2” operations, such as network topology updates, automatic monitoring of deployed AWS resources and recovery in case of their corruption / accidental removal, upgrade of Mediant VE software, etc.

Stack Manager is not involved in call processing or any other service provided by the Mediant VE. Therefore, if you are not interested in “day 2” operations provided by it, you may stop the corresponding virtual machine after Mediant VE stack deployment or even terminate it. Stack Manager uses dynamically generated Cloud Formation templates for stack deployment on AWS platform. Therefore, if you terminated it you may later delete the corresponding Mediant VE stack via Cloud Formation screen in the AWS portal.

➤ To deploy Mediant VE via Stack Manager:

1. Install the Stack Manager tool, as described in the *Stack Manager User's Manual*, which you can download from AudioCodes website at <https://www.audiocodes.com/library/technical-documents>.
2. Log into the Stack Manager tool after deployment; the following screen appears:

Figure 6-1: Stack Manager Main Screen



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

Figure 6-2: Stack Manager: Create Stack Dialog – Step 1

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

Figure 6-3: Stack Manager: 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 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:
 - **Standalone (non-HA) deployments:** Leave at default value (**none**).
 - **HA deployments:** Select the IAM role that corresponds to the deployment topology – "single zone" or "multiple zones". See Section IAM Role for Mediant VE HA Deployment for details.
9. From the 'HA mode' drop-down list, select the deployment.
10. The **VM type** is pre-selected and automatically updated based on other parameters that you define in the Create Stack dialog. If you want to modify it, check the 'Customize' checkbox next to it and then select a value from the drop-down list.
11. For HA deployment, from the 'Deployment topology' drop-down list, select **single zone** or **multiple zones**.
12. From the 'VPC' drop-down list, select the VPC where the stack will be deployed; additional fields appear in the create dialog.

Figure 6-4: Stack Manager: Create Stack Dialog – Step 3 – “Single Zone” HA

Create new stack

VPC vpc-45f3152c (DefaultVPC) ▾

HA subnet subnet-0298567cf62690236 (ipv6-b) ▾

Main subnet subnet-507ddd2d (voip2-b) ▾

1st Additional subnet -- none -- ▾

2nd Additional subnet -- none -- ▾

Public IPs Main subnet ▾

Use private IP address for management

Admin User

Username

Create Cancel

13. Select the subnets that your stack will be connected to. The list of subnets depends on **HA mode** and **Deployment topology** that you previously selected. For example, if you selected **multiple zones** HA deployment topology, two sets of subnets are shown per Availability Zone.

Figure 6-5: Stack Manager: Create Stack Dialog – Step 3 – “Multiple Zones” HA

The screenshot shows a 'Create new stack' dialog box. At the top, it says 'Create new stack'. Below that, there are several configuration sections:

- VPC:** A dropdown menu showing 'vpc-45f3152c (DefaultVPC)' with a downward arrow.
- HA subnet:** Two dropdown menus, one for '(zone 1)' and one for '(zone 2)', both showing '-- select --' with a downward arrow.
- Main subnet:** Two dropdown menus, one for '(zone 1)' and one for '(zone 2)', both showing '-- select --' with a downward arrow.
- 1st Additional subnet:** Two dropdown menus, one for '(zone 1)' and one for '(zone 2)', both showing '-- select --' with a downward arrow.
- 2nd Additional subnet:** This section is partially visible at the bottom of the dialog.

At the bottom of the dialog, there are two buttons: a blue 'Create' button and a white 'Cancel' button with a grey border.

Figure 6-6: Stack Manager: Create Stack Dialog – Step 4

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

- Public IPs:** A dropdown menu currently showing 'Main subnet'. Below it is a checkbox labeled 'Use private IP address for management' which is unchecked.
- Admin User:** Two text input fields labeled 'Username' and 'Password'.
- Advanced:**
 - SBC version:** A dropdown menu showing '7.40A.400.023'.
 - Management ports:** A text input field containing '22/tcp,80/tcp,443/tcp'.
 - Signaling ports:** A text input field containing '5060/udp,5060/tcp,5061/tcp'.
 - Use main subnet for:** A dropdown menu showing 'all traffic (management + VoIP)'.
 - Advanced config:** A large empty text area for additional configuration.

At the bottom of the dialog are two buttons: 'Create' (highlighted in blue) and 'Cancel'.

14. From the 'Public IPs' drop-down list, select which subnets need to communicate with external equipment via public IP addresses. Stack Manager assigns Elastic IP addresses to the corresponding network interfaces.
15. 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 behavior, by checking the 'Use private IP address for management' checkbox. In this case, Stack Manager uses a private IP address to communicate with the deployed stack.
16. In the 'Username' and 'Password' fields, enter the admin user credentials that will be configured on the deployed stack. You will 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 VE's Web or CLI interface) communication between Stack Manager and the deployed Mediant VE stack is not affected.
17. From the 'SBC version' drop-down list, select the Mediant VE version that you want to deploy.
18. For the 'Management ports' and 'Signaling ports' fields, enter a list of management and signaling ports respectively that should be opened on the Mediant VE. Specified ports will be configured in the corresponding network security groups assigned to Mediant VE 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 a 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 a 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. The selection effects network security groups assigned to the Mediant VE network interface connected to the Main subnet and default SIP Interface and Media Realm created on Mediant VE.
 20. In the 'Advanced config' text box, enter advanced configuration parameters, if needed. See the next sections for a partial list of supported advanced configuration parameters. Refer to *Stack Manager User's Manual* for the complete list.
 21. Click **Create** to start stack creation.
 22. Wait until stack is created.

Figure 6-7: Stack Manager: Successful Stack Creation

The screenshot shows the Stack Manager web interface. At the top, there is a navigation bar with 'stack_mgr' and links for 'Stacks', 'Configuration', 'Logs', and 'About'. On the right, there are 'Admin', 'Logout', and a help icon. Below the navigation bar, there is a '+ Create new stack' button and a search box. A large green message box displays the following text:

```

Creating stack 'mediant-ve'
Initializing AWS client... done
Creating stack..... done
Waiting until SBC is up..... done
Creating system snapshot... done

Use http://3.74.99.83 to connect to the management interface.

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

Below the message box, there is a section titled 'Stacks' containing a table with the following data:

Name	Type	Environment	State	Alarms	IP Address
mediant-ve	Mediant VE	AWS	creating		3.74.99.83

6.1 Public IP Addresses

During Mediant VE 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.

For each assigned public (Elastic) IP address, Stack Manager creates corresponding entries in the NAT Translation configuration table (of Mediant VE), 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 attach multiple public (Elastic) IP addresses to the same network interface. This may be done by specifying the **public_ips** advanced configuration parameter in the **Advanced Config** section during stack creation, or updating the **Public IPs** parameter for existing stack via the **Modify** action.



Note: When the **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.

■ public_ips

Contains 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:

```
public_ips = main:2,additional1
```

attaches two Elastic IP addresses to the network interface connected to the Main subnet (eth0 for standalone deployment, eth1 for HA deployment) and one Elastic IP address to the network interface connected to the Additional 1 subnet (eth1 for standalone deployment, eth2 for HA deployment).

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 deployment type:

- **Standalone deployments:** 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.
- **HA 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.
- **HA deployments in multiple availability zones:** Behavior is similar to standalone deployments.

6.2 Private IP Addresses for Standalone Deployments



Note: This section applies only to **standalone deployments** (as described in Section 2.1).

Standalone deployments use primary private IP addresses per interface.

For subnets that have public (Elastic) IP address assigned (as described in Section 6.1), 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 public (Elastic) IP address assigned, communication with other equipment (inside the VPC or between “connected” VPCs) happens via the interface’s primary 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 the **additional_ips** advanced configuration parameters in the **Advanced Config** section during stack creation, or update the **Additional IPs** parameter for the existing stack via the **Modify** action.

- **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:

```
additional_ips = main,additional1:2
```

creates one additional private IP address on the network interface connected to the Main subnet (eth0 for standalone deployment, eth1 for HA deployment) and two additional private IP addresses on the network interface connected to the Additional 1 subnet (eth1 for standalone deployment, eth2 for HA deployment).

The 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:

```
HA mode: disable
Main Subnet: <main-subnet-id>
1st Additional Subnet: <additional-subnet-id>
Public IPs: Main subnet
Advanced Config:
    additional_ips = main,additional1
```

creates the following networking configuration:

- **eth0** – one primary and one secondary IP addresses:
 - primary IP address – assigned with an Elastic IP address (due to the **Public IPs** configuration parameter)
 - 1st secondary IP address – created due to the **additional_ips** advanced configuration parameter containing “main” element

- **eth1** – one primary and one secondary IP address:
 - primary IP address – first “operational” private IP address, created implicitly
 - 1st secondary IP address – created due to the **additional_ips** advanced configuration parameter containing “additional1” element

6.3 Private IP Addresses for HA Deployments in Single Availability Zone



Note: This section applies only to **HA deployments in a single availability zone** (as described in Section 2.2).

HA deployments in a single availability zone don't use primary Primary IP addresses on eth1, eth2 and eth3 interfaces (connected to Main, 1st and 2nd Additional subnets correspondingly), because they can't be moved between two Mediant VE instances during activity switchover. Instead, secondary IP addresses are created and used.

For subnets that have public (Elastic) IP address assigned – as described in Section 6.1 – Elastic IP address is mapped (by AWS) to the corresponding interface's first “operational” (secondary) private IP address and NAT Translation table is configured (on Mediant CE) to reflect this mapping.

For subnets that don't have public (Elastic) IP address assigned, communication with other equipment (inside the VPC or between “connected” VPCs) happens via the interface's first “operational” (secondary) 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 some network interface, specify **additional_ips** advanced configuration parameters in **Advanced Config** section during stack creation, or update **Additional IPs** parameter for existing stack via the **Modify** action.

- **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:

```
additional_ips = main,additional1:2
```

creates one additional private IP address on the network interface connected to the Main subnet (eth0 for standalone deployment, eth1 for HA deployment) and two additional private IP addresses on the network interface connected to the Additional 1 subnet (eth1 for standalone deployment, eth2 for HA deployment).

The 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:

```
HA mode: enable
Deployment Topology: single zone
HA Subnet: <ha-subnet-id>
Main Subnet: <main-subnet-id>
```

```

1st Additional Subnet: <additional-subnet-id>
Public IPs: Main subnet
Advanced Config:
    additional_ips = main,additional1
  
```

creates the following networking configuration:

- **eth0** – one primary IP address (used for internal traffic between Mediant VE instances and for communication with AWS EC2 API)
- **eth1** – one primary and two secondary IP addresses:
 - primary IP address is not used because it can't be moved between Mediant VE instances in case of switchover
 - 1st secondary IP address - first “operational” private IP address, created implicitly and assigned with an Elastic IP address (due to the **Public IPs** configuration parameter)
 - 2nd secondary IP address - created due to the **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 Mediant VE instances in case of switchover
 - 1st secondary IP address – first “operational” private IP address, created implicitly
 - 2nd secondary IP address – created due to the **additional_ips** advanced configuration parameter containing “additional1” element

6.4 Private IP Addresses for HA Deployments in Multiple Availability Zones



Note: This section applies only to **HA deployments in multiple availability zones** – as described in Section 0.

HA deployments in multiple availability zones use primary private IP addresses on each interface.

For subnets that have public (Elastic) IP address assigned – as described in Section 6.1 – Elastic IP address is mapped (by AWS) to the corresponding interface's primary private IP address and NAT Translation table is configured (on Mediant CE) to reflect this mapping.

For subnets that don't have public (Elastic) IP address assigned, Virtual IP addresses – as described in Section 2.3.1 – are allocated and should be used for communication inside the VPC or between VPCs connected via the Transit Gateway. Private IP addresses may appear in Interfaces table, but should not be used by applications (e.g. SIP Interfaces).

If you want to enable communication via both public (Elastic) and private (virtual) IP addresses on the same subnet, specify **additional_ips** advanced configuration parameters in **Advanced Config** section during stack creation, or update **Additional IPs** parameter for existing stack via the **Modify** action.

■ additional_ips

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

For example:

```
HA mode: enable
Deployment Topology: multi zone
Public IPs: Main subnet
Advanced Config:
    additional_ips = main
```

attaches both public (Elastic) and virtual IP address to the Main subnet (eth1).

6.5 Deployment Troubleshooting

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

If Mediant VE 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.

7 Security Groups

7.1 Default Security Groups

For Mediant VE deployed via the CloudFormation template or Stack Manager, the following security groups are automatically created:

- **OAM** – security group for management traffic
- **Signaling** – security group for SIP traffic
- **Media** – security group for RTP/RTCP traffic
- **HA** – security group for internal traffic between Mediant VE instances (applicable to HA deployments only)

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

Table 7-1: Assignment of Default Security Groups

Security Group	Subnet Names	Interface Names for HA Deployment	Interface Names for Standalone Deployment
HA	HA	eth0	n/a
OAM	Main	eth1	eth0
Signaling	Main, 1 st Additional, 2 nd Additional, ...	eth1, eth2, eth3, ...	eth0, eth1, eth2, ...
Media	Main, 1 st Additional, 2 nd Additional, ...	eth1, eth2, eth3, ...	eth0, eth1, eth2, ...

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

Table 7-2: Inbound Rules for Default Security Groups

Security Group	Traffic	Protocol	Port	Source
OAM	SSH	TCP	22	0.0.0.0/0
	HTTP	TCP	80	0.0.0.0/0
	HTTPS	TCP	443	0.0.0.0/0
Signaling	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
Media	RTP, RTCP	UDP	6000-65535	0.0.0.0/0

Security Group	Traffic	Protocol	Port	Source
HA	Internal	UDP	669	HA security group
	Internal	UDP	680	HA security group
	Internal	TCP	80	HA security group
	Internal	TCP	2442	HA security group

Outbound rules are configured by default to allow all traffic.

7.2 Adjusting Default Security Groups

The default **OAM**, **Signaling** and **Media** security groups are configured by default to accept traffic from all sources, which constitutes a significant security risk. It is 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 in the **HA** security group allow only traffic that originates from instances that are attached to this security group. Therefore, there is typically no need to modify them.

For Mediant VE deployed via Stack Manager 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.

For other deployment methods, you can update inbound and/or outbound rules of the default security groups via the AWS console, CLI, or CloudFormation interfaces.

If you update outbound rules of the **HA** security group, make sure to include the following minimal required rules:

Table 7-3: Minimal Required Outbound Rules for HA Security Group

Type	Protocol	Port Range	Destination	Description
All	All	All	HA security group	Internal traffic between Mediant VE instances
HTTP	TCP	80	169.254.169.254/32	Communication with EC2 instance meta-data service
HTTPS	TCP	443	A.B.C.D/32	Communication with EC2 API endpoint. Replace A.B.C.D with the actual 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.

7.3 Using Custom Security Groups

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

For Mediant VE deployed via Stack Manager, such configuration can be done via the following stack advanced configuration parameters:

- **ha_nsg_id**

Defines a custom security group to be used instead of the default **HA** security group.

For example:

```
ha_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 via the following stack advanced configuration parameters:

- **nsg_id_ethX**

Defines a custom security group for a specific network interface instead of default security groups. Multiple security groups can be specified via a comma-separated list.

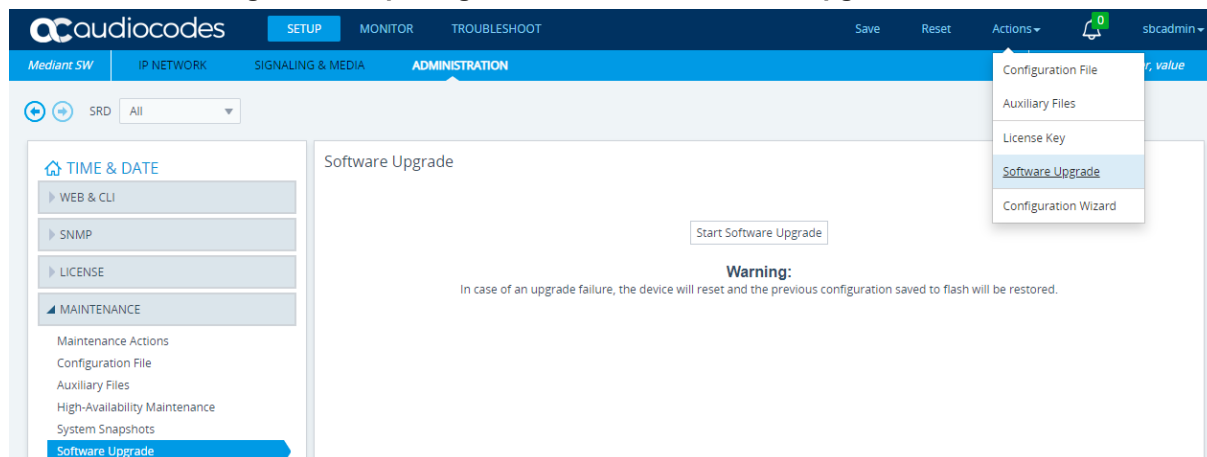
For example:

```
nsg_id_eth0 = sg-123456
nsg_id_eth1 = sg-34567,sg-56789
```

8 Upgrading the Software Version

You may upgrade the software version of the deployed Mediant VE software using the software version file (.cmp) through the Web or CLI interface. For example, open the Web interface, and then click **Action > Software Upgrade** on the toolbar to open the Software Upgrade wizard.

Figure 8-1: Opening Web Interface's Software Upgrade Wizard



Upgrading the Mediant VE using the software version file (.cmp) may be performed only within the same OS version stream.

The following streams are available:

- 7.20A stream – based on OS Version 6
- 7.20CO stream – based on OS Version 8
- 7.40A stream – based on OS Version 8

For example, if your Mediant VE is currently running Software Version 7.20A.256.396 (i.e., 7.20A stream, based on OS Version 6), you may use the 7.20A.258.010 .cmp file to upgrade it to a later version (also based on OS Version 6). However, you may not use 7.20CO.258.034 .cmp file to perform a similar upgrade to a version from the 7.20CO stream (based on OS Version 8).

If you want to upgrade Mediant VE deployed with a version from the 7.20A stream (based on OS Version 6) to a version from 7.20CO or 7.40A streams (based on OS Version 8), use one of the following methods:

- **Method 1:** Deploy a new Mediant VE instance using OS Version 8 software image, configure it, and then switch live traffic to the new instance. Refer to Section 8.1 for detailed instructions.
- **Method 2:** Rebuild the existing Mediant VE instance from the new OS Version 8 image. Refer to Section 8.2 for detailed instructions.

Advantages and disadvantages of each method are listed in the following table:

Method	Advantages	Disadvantages
Method 1	<ul style="list-style-type: none"> ■ If any problems with the new software version (based on OS Version 8) occur, live traffic may be switched back to the old instance, running the old software version. 	<ul style="list-style-type: none"> ■ Requires the use of additional AWS resources for the duration of the upgrade. ■ Requires a change of IP addresses (both public and private) and

Method	Advantages	Disadvantages
	<ul style="list-style-type: none"> Traffic may gradually be moved to a new instance (assuming that VoIP equipment that sent the traffic towards the SBC supports such functionality), thereby providing better control over the upgrade process and minimizing service downtime. 	<p>therefore, requires reconfiguration of VoIP equipment that communicates with the SBC.</p> <ul style="list-style-type: none"> Requires a new License Key for the new Mediant VE instance.
Method 2	<ul style="list-style-type: none"> Doesn't require additional AWS resources. Preserves public and private IP addresses of the deployed SBC instance. 	<ul style="list-style-type: none"> Requires a new License Key after the upgrade (because SBC serial number changes). Service is unavailable while the instance is rebuilt (typically for 5-10 minutes).

8.1 Method 1 – Side-By-Side Deployment of New Version

This section describes the upgrade of the Mediant VE instance running software version from the 7.20A stream (based on OS Version 6) to a version from the 7.20CO or 7.40A streams (based on OS Version 8) via side-by-side installation of a new Mediant VE instance and gradual migration of live traffic from the old to the new instance.

➤ To perform upgrade via "side-by-side deployment" method:

1. Deploy a new Mediant VE instance using OS Version 8 image via one of the following means:
 - For standalone Mediant VE deployment using AWS EC2 console (as described in Section 44), choose version from 7.20CO or 7.40A streams based on OS Version 8
 - For HA Mediant VE deployment using CloudFormation Service (as described in Section 0), choose the AMI ID that corresponds to version from 7.20CO or 7.40A streams based on OS Version 8
 - Using Stack Manager (as described in Section 6), choose **OS Version = 8** during the deployment.

Connect the new Mediant VE instance to the same VPC and Subnets as the existing Mediant VE instance.

2. Download the configuration file (.ini) from the existing Mediant VE instance (**Actions > Configuration File > Save INI File**).
3. Remove all networking configuration from the downloaded file, using one of the following methods:
 - Using the ini_cleanup.py script from the *Mediant VE Installation Kit*, which is available on www.audiocodes.com portal.


```
# python ini_cleanup.py old.ini new.ini
```
 - Manually: Open the file in a text editor (e.g. Notepad++), and then delete the following elements:
 - ◆ Configuration tables: PhysicalPortsTable, EtherGroupTable, DeviceTable, InterfaceTable, MtcEntities
 - ◆ Configuration parameters: HARemoteAddress, HAUnitIdName, HARemoteUnitIdName, HAPriority, HARemotePriority, HALocalMAC, HARemoteMAC

4. Load the "cleaned up" configuration file to the new Mediant VE instance as an incremental INI file (**SETUP > ADMINISTRATION > MAINTENANCE > Auxiliary Files > INI file (incremental)**).
5. Obtain, activate and apply the license to the new Mediant VE instance, as described in Section 9.1.
6. Switch live traffic from the old Mediant VE instance to the new one. This typically requires a change in the SBC's IP address in the VoIP equipment that communicates with the SBC. Consider performing gradual traffic migration if your VoIP equipment supports it. For example, switch 10% of your live traffic to the new Mediant VE instance first, verify that it is processed as expected, and only then switch the rest of the traffic.
7. After all live traffic is switched to the new Mediant VE instance and service operates normally, delete the old Mediant VE instance.

8.2 Method 2 – Rebuild Existing Mediant VE Instance from New Image

This section describes the upgrade procedure of Mediant VE instance running software version from the 7.20A stream (based on OS Version 6) to a version from the 7.20CO or 7.40A streams (based on OS Version 8) via a rebuild of existing Mediant VE instance from a new image.

The detailed procedure differs depending on the Mediant VE topology (HA or standalone) and deployment method.

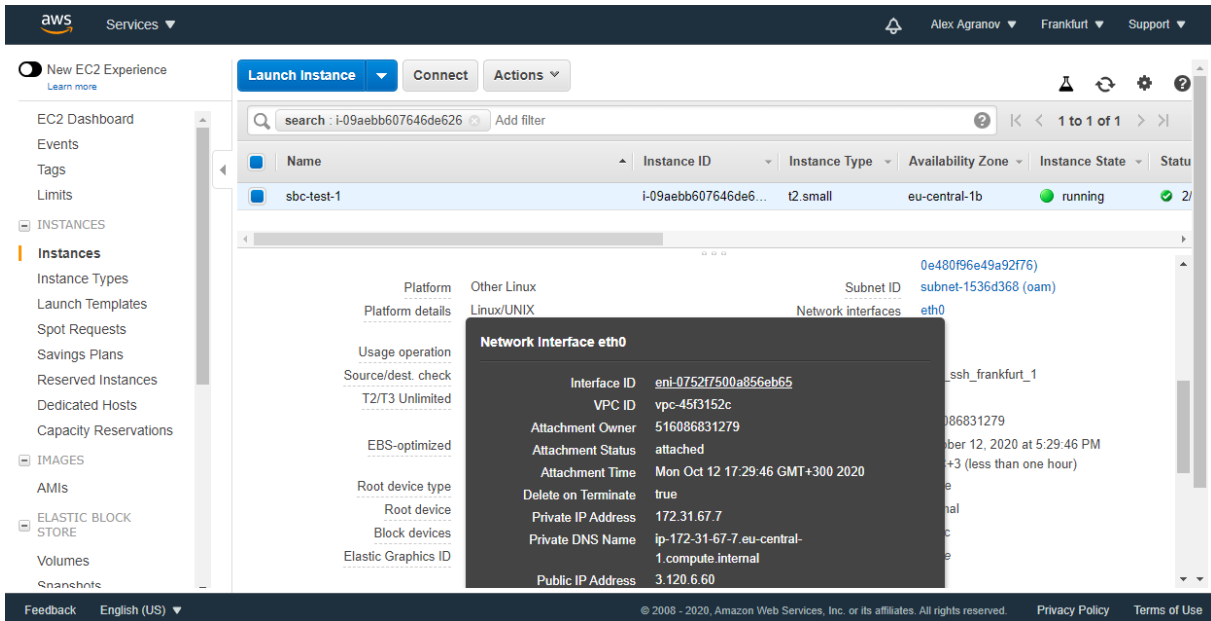
8.2.1 Rebuilding Existing Standalone Mediant VE Instance Deployed via AWS EC2 Console from New Image

The described process preserves all IP addresses (private and public) assigned to the Mediant VE instance, as well as most of the SBC configuration. However, the following configuration elements will be lost and must be manually restored afterwards:

- TLS Contexts configuration (certificates and private keys)
 - Auxiliary files (e.g., Pre-recorded Tone files)
 - License Keys (because the serial number of rebuilt instances changes)
- **To rebuild existing standalone Mediant VE instance deployed via AWS EC2 console from new image:**
1. Download the Configuration Package from the Mediant VE instance: **Actions > Configuration File > Save Configuration Package**.
 2. Open the EC2 console at <https://console.aws.amazon.com/ec2>.
 3. Navigate to the **Instances** page, and then locate your Mediant VE instance.

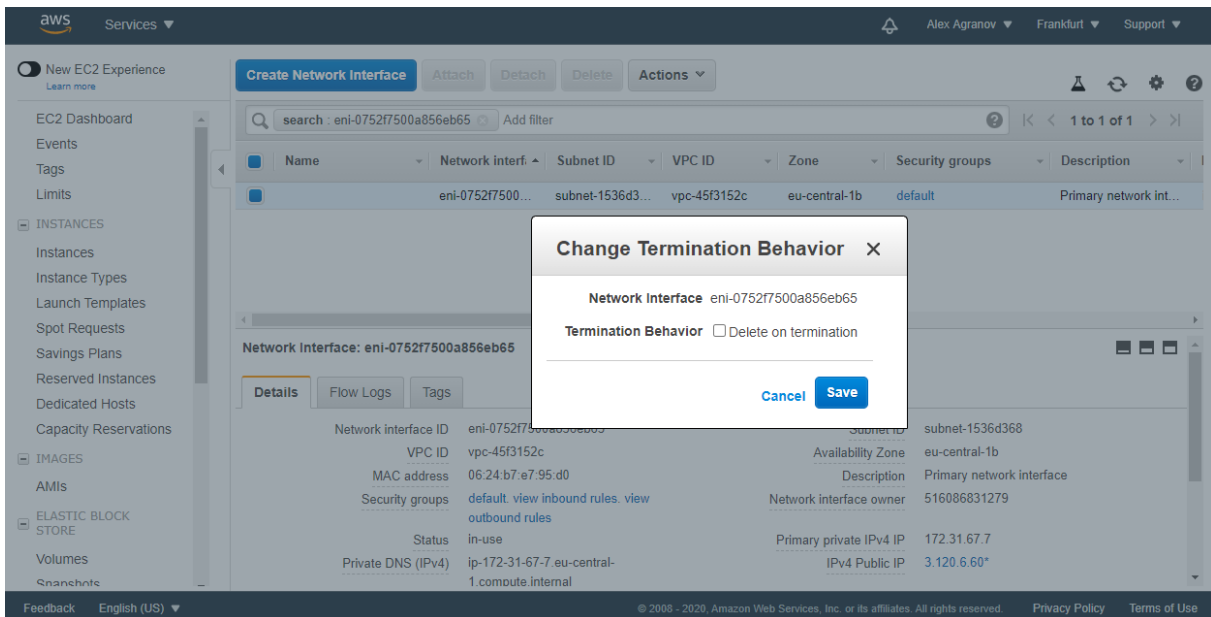
- Find network interfaces associated with your instance.

Figure 8-2: Finding Network Instances associated with EC2 Instance



- For each network interface:
 - Navigate to the specific interface in the **Network Interfaces** page.
 - Write down the interface ID (eni-xxxxxxx); you will need it in the next steps.
 - Click **Actions** > **Change Termination Behavior** and clear the 'Delete on termination' check box.

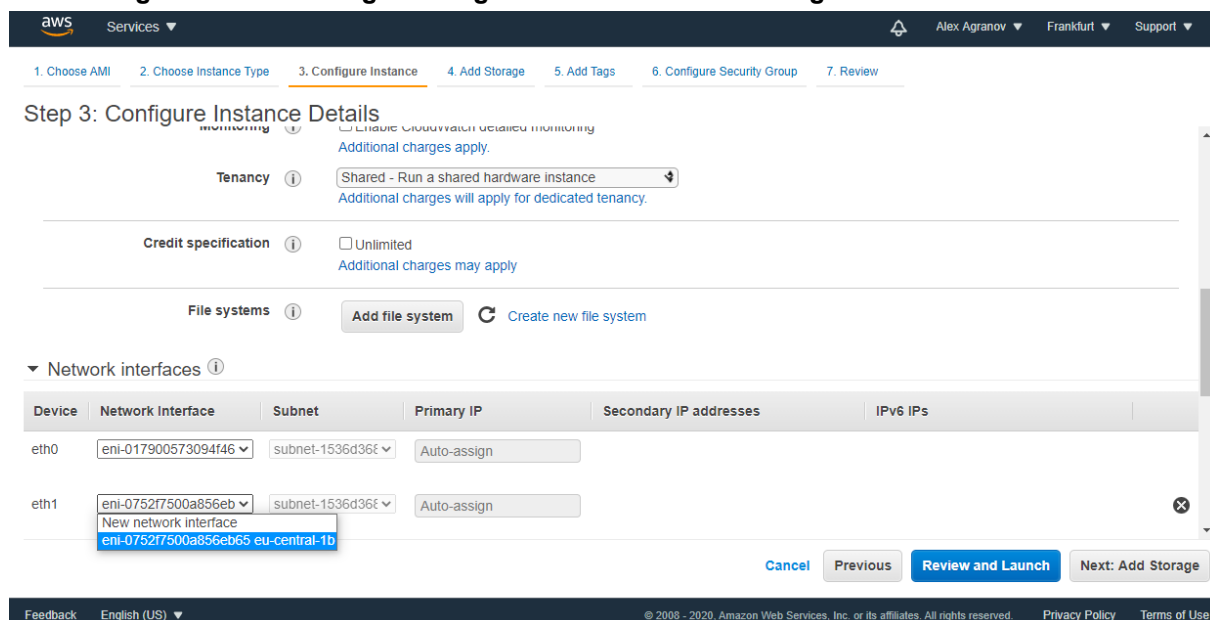
Figure 8-3: Changing Termination Behavior of Network Interface



- Navigate back to your Mediant VE instance on the **Instances** page.
- Click **Action** > **Instance State** > **Terminate** to terminate the instance. If asked if you want to release Elastic IPs, choose to preserve them.

8. Navigate to the AWS Marketplace at <https://console.aws.amazon.com/marketplace> and start a new instance deployment, as described in Section 4.
9. Choose a version from the 7.20CO or 7.40A streams, which are based on OS Version 8.
10. In the “Step 3: Configure Instance Details” screen:
 - a. Select **VPC** where the old Mediant VE instance was deployed.
 - b. Select the **Subnet** that the old Mediant VE instance’s 1st network interface was connected to.
 - c. Scroll to the bottom of the page.
 - d. Under **Network Interfaces**, select an existing network interface that was used by the old Mediant VE instance. If your instance had a second network interface, then add it and choose the corresponding existing network interface.

Figure 8-4: Choosing Existing Network Interfaces during EC2 Instance Creation



11. Proceed with new instance deployment.
12. Wait until the new Mediant VE instance is deployed and fully starts (it may take up to 5 minutes). Navigate to the **Instances** page, and then check the *instance-id* of the deployed instance.
13. Log in to the new Mediant VE instance using the following default credentials:
 - Username: **Admin**
 - Password: *instance-id*
14. Load the Configuration Package file, which was saved in Step 1, back to the device (**Actions > Configuration File > Load Configuration Package**).
15. Restore parts of the Mediant VE configuration that have been lost during the rebuild, namely, TLS Contexts configuration (certificates and private keys) and Auxiliary files.
16. Obtain, activate and apply the license to the new Mediant VE instance, as described in Section 9.

Your Mediant VE is now running the new software version based on OS Version 8 and is fully operational.

8.2.2 Rebuilding Existing High-Availability (HA) Mediant VE Deployed via AWS EC2 Console from New Image

Rebuilding of the existing High-Availability (HA) Mediant VE deployed via AWS EC2 Console using CloudFormation template consists of the following steps:

1. Updating stack with change set #1 that deletes all EC2 instances and related resources.
2. Updating stack with change set #2 that restores all EC2 instances and related resources, using a new image ID (with the new Mediant VE version based on OS Version 8).

The described process preserves all IP addresses (private and public) assigned to the Mediant VE instance, as well as most of the SBC configuration. However, the following configuration elements will be lost and must be manually restored after it:

- TLS Contexts configuration (certificates and private keys)
- Auxiliary files (e.g., Pre-recorded Tone files)
- License Keys (because the serial number of rebuilt instances changes)

➤ **To rebuild existing high-availability (HA) Mediant VE deployed via AWS EC2 console from new image:**

1. Make sure that the 1st SBC instance (SBC-1) is currently active. If not, perform a switchover to make it active.

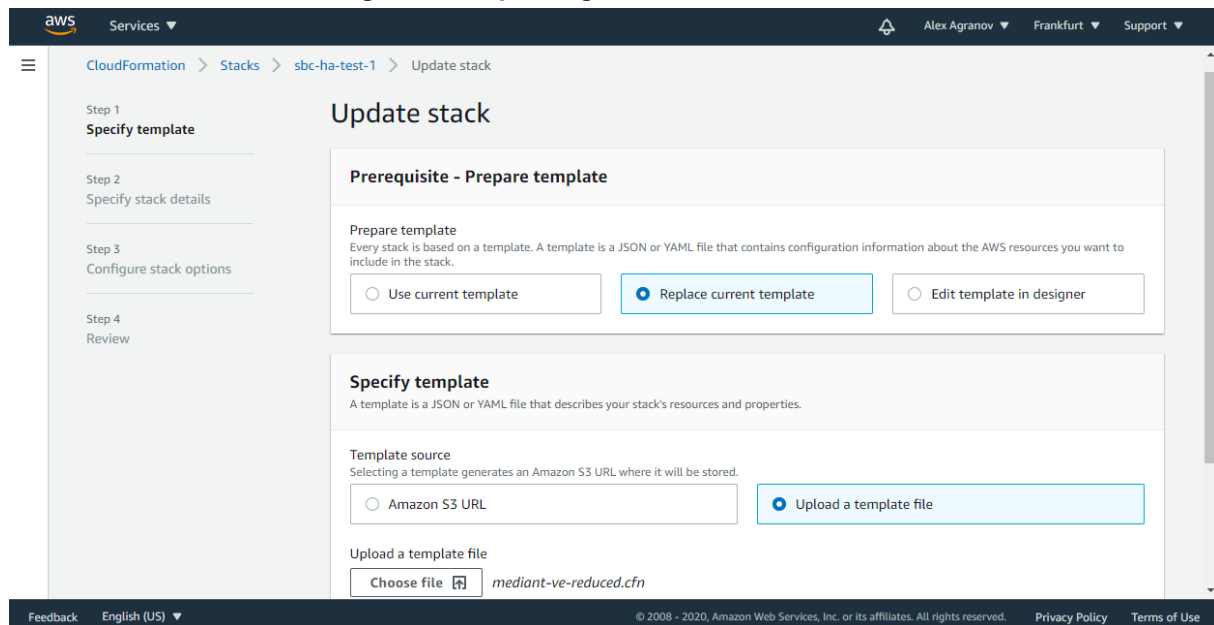


Note: Secondary IP addresses move during activity switchover. If the 2nd SBC instance is currently active, secondary IP addresses are assigned to it and therefore, stack runtime configuration doesn't match the CloudFormation template. This will result in a failure in the stack update procedure, described below.

2. Download the Configuration Package from the Mediant VE instance (**Actions > Configuration File > Save Configuration Package**).
3. Open the CloudFormation console at <https://console.aws.amazon.com/cloudformation>.
4. Locate the Mediant VE stack.
5. Switch to the **Template** tab, copy the current stack template to the clipboard and paste it into a new file on your PC. Name the file "mediant-ve.cfn".
6. Create a copy of the file "mediant-ve.cfn" and name it "mediant-ve-reduced.cfn". Edit the copied file "mediant-ve-reduced.cfn" as follows:
 - a. Remove the following elements from the **Resources** array:
 - sbc1
 - sbc2
 - sbc1eth2Attachment
 - sbc2eth2Attachment
 - sbc1eth3Attachment
 - sbc2eth3Attachment
 - eth1EIPAssociation
 - recoveryTestAlarmSbc1
 - recoveryTestAlarmSbc2

- b. Remove the following elements from the **Outputs** array:
 - sbc1Instanceid
 - sbc2Instanceid
7. In the CloudFormation screen, click **Update**.
8. Choose **Replace current template**, upload the “mediant-ve-reduced.cfn” file from your PC, and then click **Next**.

Figure 8-5: Updating Cloud Formation stack



9. In the subsequent screens, click **Next** to accept default parameters, and then click **Update stack**.
10. While the stack is updated, its state changes to "UPDATE_IN_PROGRESS". Wait until the update is complete and the stack state changes to "UPDATE_COMPLETE".
11. In the CloudFormation screen, click **Update** again.
12. Choose **Replace current template**, upload the “mediant-ve.cfn” file from your PC, and then click **Next**.
13. In the **Specify stack details** screen, modify the **Amazon Machine Image (AMI)** parameter to the value of the AMI that corresponds to a new Mediant VE version (based on OS Version 8). Use AWS Marketplace <https://console.aws.amazon.com/marketplace> to determine the correct AMI ID that corresponds to the region where Mediant VE is deployed.
14. In the subsequent screens, click **Next** to accept default parameters, and then click **Update stack**.
15. While the stack is updated, its state changes to "UPDATE_IN_PROGRESS". Wait until the update is complete and the stack state changes to "UPDATE_COMPLETE".
16. Log in to the new Mediant VE instance using the default admin credentials – **adminUsername** and **adminPassword** – listed in the **Outputs** tab.



Note: If you copy/paste the *instance-id* from the **Outputs** tab, your browser may append a space to the copied value, thus making it invalid. Therefore, it is recommended to type the *instance-id* manually.

17. Load the Configuration Package file, which you saved in Step 1, back to the device (**Actions > Configuration File > Load Configuration Package**).
18. Restore parts of the Mediant VE configuration that have been lost during the rebuild,

namely, TLS Contexts configuration (certificates / private keys) and Auxiliary files.

19. Obtain, activate and apply the license to the new Mediant VE instance, as described in Section 9.

Your Mediant VE is now running the new software version based OS Version 8 and is fully operational.

8.2.3 Rebuilding Existing Mediant VE Deployed via Stack Manager

This chapter describes the upgrade of Mediant VE instance running software version from 7.20A stream (based on OS Version 6) to a version from 7.20CO or 7.40A streams (based on OS Version 8) via a rebuild of existing Mediant VE instance from a new image using the Stack Manager.

The described process preserves all IP addresses (private and public) assigned to the Mediant VE instance, as well as most of the SBC configuration. However, the following configuration elements will be lost and must be manually restored afterwards:


- TLS Contexts configuration (certificates and private keys)
- Auxiliary files (e.g., Pre-recorded Tone files)
- License Keys (because the serial number of rebuilt instances changes)

➤ **To rebuild existing Mediant VE deployed via Stack Manager:**

1. Connect to the Stack Manager Web interface.
2. Click the corresponding stack name.
3. Click **Modify**, and then change the **OS Version to 8**.
4. Click **Update** to rebuild the stack.
5. Wait for the **Update** operation to complete. The operation typically takes 10-15 minutes, during which all VM instances are rebuilt and service is unavailable. Mediant VE configuration, including private and public IP addresses is preserved.
6. Restore parts of the Mediant VE configuration that have been lost during the rebuild, namely, TLS Contexts configuration (certificates / private keys) and Auxiliary files.
7. Obtain, activate and apply the license to the signaling components, as described in Section 9.

Your Mediant VE is now running the new software version based on OS Version 8 and is fully operational.

Figure 8-6: Upgrading Mediant VE to New Image Based on OS Version 8 via Stack Manager



9 Licensing the Product

Mediant VE SBC is available in AWS Marketplace as two different products:

- **Mediant VE Session Border Controller (SBC):** This product includes a trial license (see below) and requires purchase of a production license from AudioCodes.
- **Mediant VE Session Border Controller (SBC) – PAYG:** This product includes a pay-as-you-go license that enables Customers to use the SBC as much as needed and pay for the actual service consumed via their AWS account billing.

If you installed the regular (not pay-as-you-go) version of the Mediant VE SBC product, your product includes a trial license that includes the following:

- Three concurrent sessions (signaling and media).
- Three user registrations (far-end users).
- Transcoding capabilities – in order to activate them you need to configure the 'SBC Performance Profile' parameter to **Optimize for Transcoding** (for more information, refer to the *User's Manual*).

Once you are finished evaluating the product you need to obtain, activate and then install the production SBC license.

9.1 Obtaining and Activating a Purchased License Key



Note: This and the following sections are not applicable to **Mediant VE SBC – PAYG** product, which doesn't require any additional license.

For the product to provide you with all your capacity and feature requirements, you need to purchase a new License Key that allows these capabilities. The following procedure describes how to obtain and activate your purchased License Key.



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 HA, each unit has its own Serial Number, Product Key and License Key. Therefore, the instructions in this section must be done for each unit.


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

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

Figure 9-1: Software License Activation Tool

2. Enter the following information:
 - **Product Key:** The Product Key identifies your specific Mediant VE SBC 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 9-2: Product Key in Order Confirmation E-mail



Note: For 1+1 High-Availability orders, you are provided with two Product Keys, one for each unit. In such cases, you need to perform license activation twice in order to obtain License Keys for both units.

- **Fingerprint:** The fingerprint is the Mediant VE SBC'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 **Send** to submit your license activation request.

- 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 Mediant VE SBC.



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

9.2 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 HA contains two License Keys - one for the active device and one for the redundant device. Each License Key has a different serial number ("S/N"), which reflects the serial number of each device in the HA system.

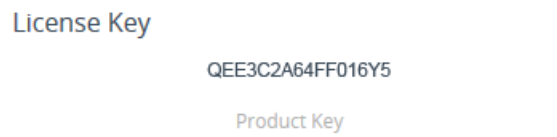
9.3 Product Key

The Product Key identifies a specific purchase of your device installation 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 > **License** folder > **License Key**). The Product Key is displayed in the read-only 'Product Key' field, as shown in the example below:

Figure 9-3: 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 9-4: Empty Product Key Field



3. Click "empty"; the following appears:

Figure 9-5: Entering Product Key



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

This page is intentionally left blank.

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>

©2026 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-11035

