Check Point
SOFTWARE TECHNOLOGIES LTD.

25 October 2020

CLOUDGUARD IAAS HIGH AVAILABILITY FOR AZURE

R80.10 AND HIGHER

Deployment Guide

Important Information

Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.

Certifications

For third party independent certification of Check Point products, see the Check Point Certifications page.

Check Point CloudGuard IaaS High Availability for Azure R80.10 and Higher Deployment Guide

Latest Version of this Document in English

Open the latest version of this document in a Web browser.
Download the latest version of this document in PDF format.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.
Please help us by sending your comments.

Revision History

Date	Description
23 September 2020	Updated steps in *"Upgrading a Check Point CloudGuard IaaS High Availability Solution to a Newer Version" on page 39*
05 April 2020	Updated *"Changing Template Components" on page 36*.
20 February 2020	Added links for getting the latest version of the test script in *"Testing and Troubleshooting" on page 30*.
05 February 2020	Added Note- By default, every Check Point Security Gateway and Security Management Server's Gaia Portal is accessible from the internet by browsing to `http://<virtual-machine-public-ip>`. Restricting access to the Gaia Portal is possible by configuring a Network Security Group, or by configuring the settings on the Check Point Security Gateway and Security Management Server.
12 January 2020	In *"Workflow for Setting Up a High Availability Cluster in Azure" on page 15*, updated .*"Step 5: Configure Cluster Objects in SmartConsole" on page 22*
31 December 2019	In *"Additional Information" on page 30*, updated steps: "To get the latest version of the test script".
10 November 2019	Updated step 2 in *"Upgrading a Check Point CloudGuard IaaS High Availability Solution to a Newer Version" on page 39*
25 August 2010	First release of this document.

Table of Contents

CloudGuard IaaS High Availability for Azure

Check Point and Microsoft have partnered to deliver a best-in-class experience for customers looking to extend advanced security protections to their Azure public and hybrid environments. Seamlessly integrating with the Azure and Azure Stack cloud infrastructures, CloudGuard for Microsoft Azure provides reliable and secure connectivity to public cloud assets while protecting applications and data with industry-leading threat prevention. Additionally, CloudGuard helps organizations by dramatically simplifying security management and policy enforcement across private, hybrid, and public cloud networks.

IT organizations can now achieve an advanced security posture that moves with Virtual Applications as they migrate from data centers to Azure hybrid cloud environment. As an Azure certified technology solution, CloudGuard compliments Azure cloud security controls to enable you to easily and seamlessly secure your assets in the cloud with elastic scalability and high availability using a cloud security solution integrated with both Azure and Azure Stack.

Prerequisites

To set up your system most efficiently, you must be familiar with these topics:

Vendor	Topics
Microsoft Azure	■ Virtual Networks ■ Virtual Machines ■ Load Balancers ■ High Availability ports ■ Public IP addresses ■ User Defined Rules (UDR) ■ Role Based Access Control (RBAC)
Check Point	■ R80.10 and Higher ■ Check Point with Microsoft Azure

Setting Up Check Point Clusters in Azure

A cluster is a group of Virtual Machines that work together in High Availability Mode. One Cluster Member is the Active, and the second Cluster Member is the Standby. The cluster fails over from the Active Cluster Member to the Standby Cluster Member when necessary.

- Cluster Members communicate to each other with unicast IP addresses.
- For inbound, outbound, and East-West traffic, Cluster Members rely on Azure Load Balancers to represent their external and internal Virtual IP addresses. Load Balancers only forward traffic to the Active Cluster Member.

- For VPN traffic, Load Balancers use API calls to Azure to communicate the failover from the Active Cluster Member. The Standby Cluster Member then promotes itself to Active.

  During cluster failover, the Standby Cluster Member associates the private and public cluster IP addresses of the Active Cluster Member with its external interface.

Azure API authentication

To make API calls to Azure automatically, Cluster Members need Azure Active Directory credentials. Use the Role-Based Access Control (RBAC) to enable Active Directory.

The Check Point Security Management Server in the Azure Cloud, or on-premises, manages the Check Point Cluster Members

Azure Internal Load Balancer

The Internal Load Balancer deploys by default as part of the solution template. It is automatically configured to listen and forward any TCP or UDP traffic on its High Availability ports. The Internal Load Balancer gets an automatically assigned name:

```
backend-lb.
```

Azure sends probes from the source IP address 168.63.129.16 to TCP port 8117 to monitor the health of the Check Point CloudGuard IaaS Security Gateways.

Network

Follow this network diagram to configure your system. Make sure to replace the IP addresses in the sample environment with the IP addresses in your environment.

Network Diagram

The diagram.

See the routing tables below the diagram.

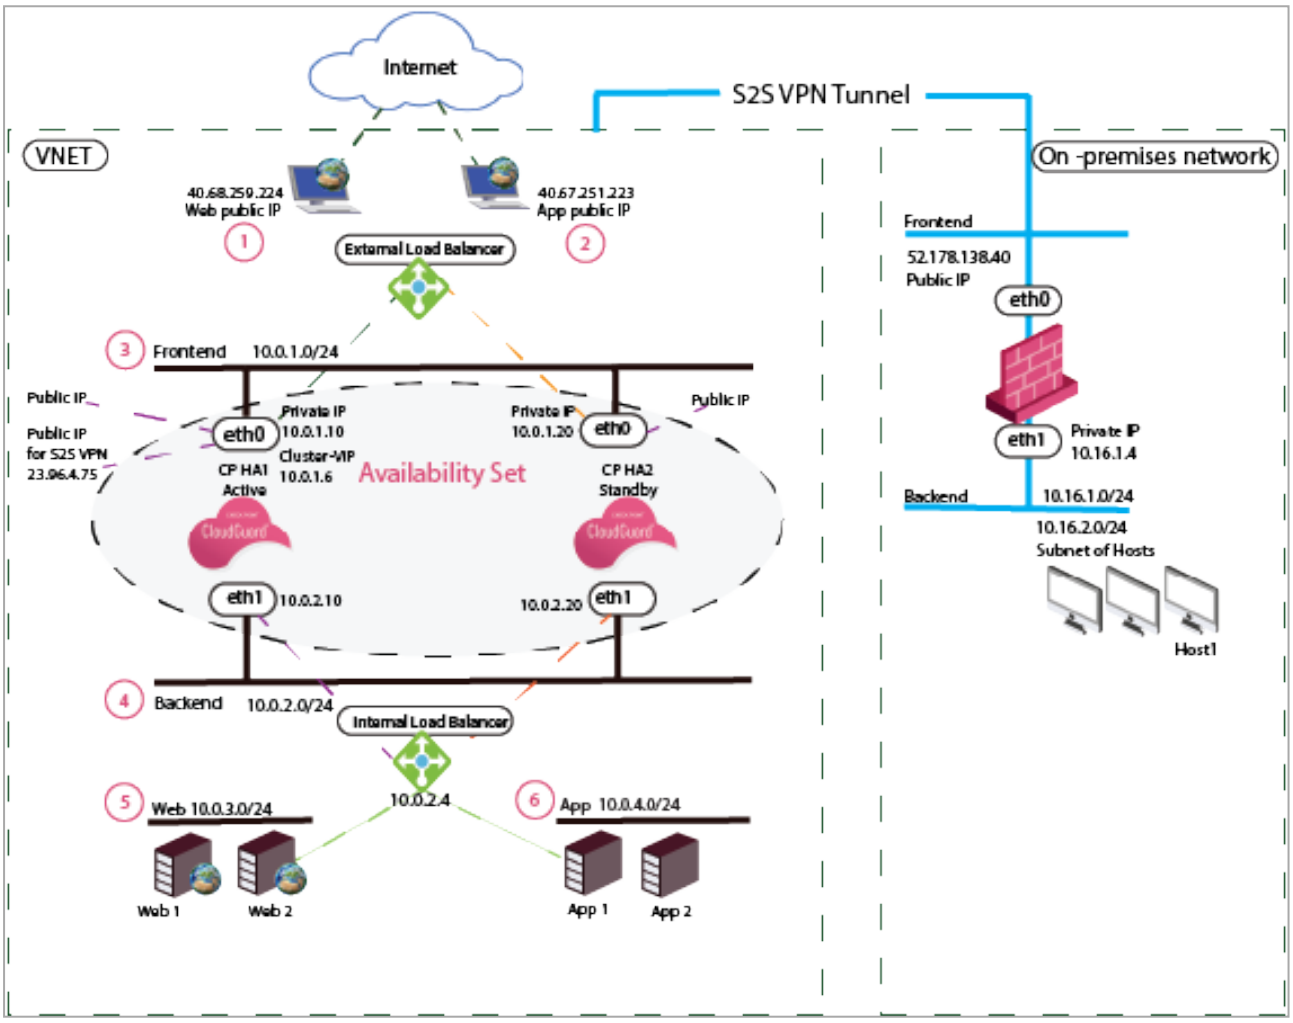

Load Balancing Rules of the External Load Balancer

1	**Example 1**	**Frontend** Web:443	**Backend port** 8081
2	**Example 2**	**Frontend** App:80	**Backend port** 8083

Frontend Routing Table - User Defined Routes (UDR)

3	**Destination**	**Nexthop**
	10.0.0.0/16	None (Drop)
	10.0.1.0/24	Virtual Network

Backend Routing Table - User Defined Routes (UDR)

4	**Destination**	**Nexthop**
	0.0.0.0/0	None (Drop)

Routing Table for Web and App - User Defined Routes (UDR)

Web and App routing tables have the same Virtual Network address, but different subnet addresses.

Web:

5	**Frontend**	**Nexthop**
	10.0.0.0/16 - *Virtual Network address*	**10.0.2.4 -***IP of the Internal Load Balancer*
	0.0.0.0/0	**10.0.2.4 -***IP of the Internal Load Balancer*
	10.0.3.0/24 (Web) - *Subnet address*	Virtual Network

App

6	**Frontend**	**Nexthop**
	10.0.0.0/16 -*Virtual Network address*	**10.0.2.4 -***IP of the Internal Load Balancer*
	0.0.0.0/0	**10.0.2.4 -***IP of the Internal Load Balancer*
	10.0.4.0/24 (App) - *Subnet address*	Virtual Network

Diagram Components

The diagram shows:

- Virtual Network in Azure that is divided into four subnets
  - Frontend
  - Backend

    - Web
    - App
- On-premises network with these components
    - Security Gateway
    - Hosts

Check Point High Availability consists of two Cluster Members, Member 1 and Member 2. Each Cluster Member has two interfaces.

When the Cluster Members are in the same Availability Set, it guarantees that the two Cluster Members are in separate fault domains. For more information, see *Manage the availability of Windows virtual machines in Azure*.

In the diagram:

- The cluster protects two web applications.
- There is Site-to-Site VPN connectivity between the Cluster Members and on-premises Security Gateways.

Each web application has:

- Public IP address
- Web server
- Application server

You must manually configure these components:

- Backend hosts
- Subnets
- Routing tables for Web and App servers

Static IP Addresses

Name	Attached to	Use
Cluster public address	The external interface of the Active Cluster Member.	VPN
Cluster private address	The external interface of the Active Cluster Member.	VPN
Member 1 public address	The external interface of Member 1.	■ External management of Member 1 ■ Internet and Azure API access Do not disable or delete this resource.

Name	Attached to	Use
Member 2 public address	The external interface of Member 2.	■ External management of Member 2 ■ Internet and Azure API access Do not disable or delete this resource.
Web	Azure Load Balancer	Public service Web
App	Azure Load Balancer	Public service App

Using the Azure Load Balancer rules to forward traffic that comes from the Internet

Note

You cannot use these ports:

- 80
- 443
- 444
- 8082
- 8880
- 8117

Azure Load Balancer rules

Frontend IP address	Frontend TCP ports	Destination IP address	Destination port
Web	HTTPS	Active Cluster Member	8081
App	HTTP	Active Cluster Member	8083

Failover

This is what happens during cluster failover:

1. The Cluster Member that fails, immediately stops responding to the Load Balancer health probes.
2. The Cluster Member that gets promoted to Active, starts responding to the Load Balancer health probes.
3. The Azure External Load Balancer and Internal Load Balancer detect the new health status of each Cluster Member, and forward traffic to the healthy Cluster Member. For more information, see *Azure Load Balancer health probes*.

   This usually happens in less than 15 seconds based on the health probe Load Balancer configuration. This affects inbound and East-West traffic inspection.
4. The Cluster Member that gets promoted to Active, uses the Azure API to associate itself with the cluster private and public IP addresses.

   This usually happens in less than 2 minutes. This affects VPN tunnel failover.

The expected failover times based on use case

Use Case	Expected Failover Time	Comments
Site-to-Site VPN	Less than 2 minutes	Depends on the Azure API.

Use Case	Expected Failover Time	Comments
Inbound inspection through the External Load Balancer	Less than 15 seconds	Depends on the Load Balancer health probe.
Outbound inspection through the Internal Load Balancer	Less than 2 minutes	Depends on the Load Balancer health probe and Azure API.
East-West inspection through the Internal Load Balancer	Less than 15 seconds	Depends on the Load Balancer health probe.

Traffic Flows

If the Management Server is in the Virtual Network, make sure to have specific routes to allow traffic between the Management Server Virtual Machine and the Cluster Members.

Note - No other Virtual Machines can be deployed in the Check Point solution subnets.

Inbound Traffic

- Traffic travels into the External Load Balancer.
- The External Load Balancer forwards the traffic to the Active Cluster Member.
- The Active Cluster Member inspects the traffic, and forwards it to the destination.

Inbound Traffic Reply

1. The traffic travels from the Web Server to the Internal Load Balancer.
2. The Internal Load Balancer forwards it to the Active Cluster Member.
3. The Active Cluster Member forwards it to the destination.

Inbound VPN Traffic

1. Packet enters the frontend NIC of the Active Cluster Member.
2. The Active Cluster Member decrypts the packet.
3. The Active Cluster Member forwards the packet to its destination.

Outbound Traffic

1. Traffic travels to an Internal Load Balancer based on the UDR.
2. The Internal Load Balancer forwards the traffic to the Active Cluster Member.
3. The Active Cluster Member inspects the traffic and forwards it to the destination.

East-West Traffic

1. Traffic travels from one of the internal servers to the Internal Load Balancer of the Check Point solution.
2. The Internal Load Balancer forwards the traffic to the Active Cluster Member.
3. The Active Cluster Member forwards the traffic to the destination.

Note:

The Internal Load Balancer deploys by default as part of the solution template and is automatically configured. It is configured to listen and forward any TCP or UDP traffic High Availability ports. It gets an automatically assigned name: `backend-lb`.

Azure sends probes from the source IP address 168.63.129.16 to TCP port 8117 to monitor the health of the Cluster Members.

Intra-Subnet Traffic

Traffic travels freely in the subnet without inspection.

Workflow for Setting Up a High Availability Cluster in Azure

Step 1: Deploy with a Template in Azure

Deploy this solution through the Azure Portal. If you use a different environment than the Standard Azure environment, see *Using a Different Azure Cloud Environment*.

- To access the Standard Azure environment, from the Azure Marketplace, see the Azure standard portal.
- To access the Azure US Government environment, from the Azure Marketplace, see the Azure US Government portal.

Notes:

- Standard Load Balancers and High Availability ports are not available on the Azure Government Cloud environment.
- By default, every Check Point Security Gateway and Security Management Server's WebUI is accessible from the internet by browsing to `http://<virtual-machine-public-ip>`. Restricting access to the WebUI is possible by configuring a Network Security Group, or by configuring the Check Point Gateway and Management Server settings.

When the template shows, enter information for these parameters:

Parameter	Description
Cluster object name	Name of the cluster object resource group.
Credentials	Public key or user name and password for SSH connections to the Cluster Members.
Subscription	Azure subscription into which the cluster object is deployed.
Resource group	Azure resource group into which the cluster object is deployed.
Location	Location into which the cluster object is deployed.
License	Type of license: ■ Bring your own license (BYOL) ■ Pay as you go (PAYG)
Virtual Machine size	Size of each Virtual Machine instance in the cluster object.
SIC	SIC key to the Security Management Server.

Parameter	Description
Network setting	1. Pre-existing Virtual Network and its subnets 2. Name of a new Virtual Network and subnets, into which the cluster object is deployed. **Notes:** When you use pre-existing subnets, make sure that: ■ No other Virtual Machines are deployed in those subnets. ■ Define UDRs properly for each subnet. See *Step 3: Set Up Internal Subnets and Route Tables*. ■ A Network Security Group (NSG) is associated with your Frontend subnet to connect the External Load Balancer and Cluster Member.
Availability Zones	Use **Availability Set** (default) or **Azure Availability Zones** for your High Availability. ■ First Cluster Member is be deployed in zone 1. ■ Second Cluster Member is deployed in zone 2. **Notes:** ■ Only available if you deploy in a supported Azure location. ■ Support for Azure Availability Zones is available with template version 20190303 and above.

Components of the Check Point Solution

The Check Point deployed solution has these components:

- Frontend subnet

  The NSG is associated with the frontend subnet and allows all inbound and outbound TCP and UDP traffic.
- Backend subnet
- Two Virtual Machines configured as a Check Point cluster
- Internal Load Balancer
- External Load Balancer
- Public IP address for each Cluster Member

No other Virtual Machines can be deployed in the solution's subnet.

Notes about the template:

- You can create a new Virtual Network, or deploy into an existing Virtual Network.
- Web and App subnets are not deployed automatically.
- It does not deploy any other Virtual Machines in the solution's frontend and backend subnets.

- Virtual Machines that are launched in the backend subnets, may need Internet access to finalize provisioning. Launch these Virtual Machines only after you have applied *Hide NAT* rules on the cluster object to support this type of connectivity.
- The Check Point First Time Configuration Wizard automatically deploys after you have set up the cluster object. The cluster object is configured based on the parameters you apply.
- After the First Time Configuration Wizard completes, the Virtual Machines automatically reboot.

Important - If you deploy the solution to an existing Virtual Network, confirm that an NSG is associated with the frontend subnet that allows all inbound and outbound TCP and UDP traffic. An NSG is necessary to connect to Cluster Members successfully.

Step 2: Set Credentials in Azure

By default, the automatic service principal is deployed. If you want to create your own service principal, make sure you set credentials and assign privileges to necessary resources. Managed service identity for Virtual Machines is only available in the Azure Cloud environment.

If you deploy in other environments, you have to create your own service principal manually. See *"Creating Your Own Service Principal" below*.

Azure Credentials and the Automatic Service Principal

The Check Point cluster template automatically creates a service principal for each Virtual Machine, and assigns a Contributor role to the cluster resource group. Therefore, you do not need to create a service principal, assign it a role, and attach it to each of your individual cluster resources. For more information, see *What is managed identities for Azure resources?*

After you deploy a Check Point cluster, the automatic credentials can be found in **Azure Portal** > **Resource groups** > **cluster_resource_group** > **Access control (IAM)**. There are two service principals for each Cluster Member, each with a Contributor role.

Notes:

- If you delete the Cluster Member's Virtual Machine, the credentials are also deleted.
- Service principals never expire.

Creating Your Own Service Principal

See *How to: Use the portal to create an Azure AD application and service principal that can access resources*.

Use these parameters:

Field	Parameter
Name	*<Application_Name>* Example: `check-point-cluster`

Field	Parameter
Application type	Web-App / API

Sign-on URL	`https://localhost/<Application_Name>`
	Example:
	`https:/localhost/check-point-cluster`

After you create the application, write down these values:

- Application ID

  `client_id`
- Key Value

  `client_secret`
- Tenant ID (Directory ID)

  `tenant`

Best Practice - We recommend that you set the key to never expire. Go to your resource.

We recommend that you set the key to never expire. Go to your resource.

To create a service principal:

Step	Description
1	Click **Access control (IAM)** > **Add**.
2	Select your role.
3	Select your AD application.
4	Click **Save**.

5	Set the `client_id` and `client_secret` on each of the Cluster Members.
	From Expert Mode, run this command on each Cluster Member:
	`# azure-ha-conf --client-id '<ApplicationId>' --client-secret '<Key Value>' --force`
	Example:
	`# azure-ha-conf --client-id '5c1896fe-26b6-4a5b-8c81-34ae07c09a24' --client-secret '2G6E_|]Y&|@Il(L}-O>g' --force`
	Note - Use single quotes to avoid shell expansion.

Step	Description
6	Make sure the file syntax is correct. From Expert Mode, run this command on each Cluster Member: `# python -m json.tool $FWDIR/conf/azure-ha.json`
7	Reload the cluster Azure configuration. From Expert Mode, run this command on each Cluster Member: `# $FWDIR/scripts/azure_ha_cli.py reconf`

To revert to your previous automatic credentials:

Step	Description
1	Remove your service principal.
2	From Expert Mode, run this command on each Cluster Member: `# azure-ha-conf --system-assigned --force`
3	Assign the two service principals to each resource and to Cluster Members. For more information, see *"Components of the Check Point Solution" on page 16* > *Notes about the template*.
4	The service principal deploys automatically. If you want to create a new service principal, assign the privileges to the necessary resources and to Cluster Members. For more information, see *Manage access using RBAC and the Azure portal*.

Step 3: Set Up Internal Subnets and Route Tables

You can use the Azure portal or the CLI to add internal subnets. You will now add the Web and App subnets to your Virtual Network.

For each internal subnet, you have to create an Azure routing table with these UDRs:

Web Route Table

#	Name	Address prefix	Nexthop-type	Nexthop-address
1	`<web-subnet>-local`	`<10.0.3.0/24>`	Virtual Network	-
2	`web-subnet-to-other-subnets`	`10.0.0.0/16`	Virtual appliance	ILB-internal-address 10.0.2.4

#	Name	Address prefix	Nexthop-type	Nexthop-address
3	`web-subnet-default`	`0.0.0.0/0`	Virtual appliance	ILB-internal-address 10.0.2.4

App Route Table:

#	Name	Address prefix	Nexthop-type	Nexthop-address
1	`<app-subnet>-local`	`10.0.4.0/24`	Virtual Network	-
2	`<app-subnet-to-other-subnets>`	`<10.0.0.0/16>`	Virtual appliance	ILB-internal-address 10.0.2.4
3	`app-subnet-default`	`0.0.0.0/0`	Virtual appliance	ILB-internal-address 10.0.2.4

Note:

If traffic inspection is needed inside the Web/App subnets, override Rule 1 in the route tables, "`<web-subnet>-local`", and "`<app-subnet>-local`".

Important - Associate the newly created routing table with the subnet to which it belongs.

If the subnet houses the Security Management Server that manages the Cluster Members, add these routes below as well. This allows the Security Management Server to communicate directly with each Cluster Member, without passing through the Active Cluster Member.

For example:

Name	Address prefix	Nexthop-type	Nexthop-address
`subnet-name-cluster_member1-management`	`cluster_member1-internal-address/32 <10.0.2.10/32>`	Virtual appliance	`cluster_member1-internal address <10.0.2.10>`
`subnet-name-cluster_member2-management`	`cluster_member2-internal-address/32 <10.0.2.20/32>`	Virtual appliance	`cluster_member2-internal address <10.0.2.20>`

Step 4: Set Up Routes on Cluster Members to the Internal Subnets

To set up route routs on Cluster Members to the internal subnets:

Step	Description
1	Connect over SSH to each of the Cluster Members.
2	Log in to Gaia Clish, or Expert mode.

3	Add this route:
	■ In Gaia Clish, run these two commands:
	`set static-route <Virtual-Network-IP-address/Prefix> nexthop gateway address <eth1-router-IP-address> on`
	`save config`
	■ In Expert mode, run this command:
	`clish -c 'set static-route <Virtual-Network-IP-address/Prefix> nexthop gateway address <eth1-router-IP-address> on' -s`
	Example:
	`set static-route 10.0.0.0/16 nexthop gateway address 10.0.2.1 on`

Parameters:

Parameter	Description
<Virtual-Network-IP-address/Prefix>	Specifies the prefix of the entire Virtual Network. Example: 10.0.0.0/16
<eth1-router-IP-address>	Specifies the first unicast IP address on the subnet, to which the `eth1` is connected. Example: 10.0.2.1

Notes:

- If the Virtual Network comprises several non-contiguous address prefixes, repeat the command for each prefix.
- For vNET Peering:
  - Add a compatible route on each peer network.
  - Add the route for vNET peering to each Cluster Member.

Step 5: Configure Cluster Objects in SmartConsole

To configure Cluster objects in SmartConsole

Step	Description
1	Click the **Objects** menu > **More object types** > **Network Object** > **Gateways and Servers** > **Cluster** > **New Cluster**.
2	Select **Classic Mode**. The **Gateway Cluster Properties** window opens.
3	Enter a **Name**. Example: `checkpoint-cluster`
4	In the **IPv4 Address** field, enter the public address allocated for the cluster. **Note** - You can find the cluster IP address in the Azure portal when you select the Active Cluster Member's primary **NIC** > **IP configuration** > **"cluster-vip"**.

Step	Description
5	Select the **Cluster Members** tab. a. Click **Add** > **New Cluster Member**. b. In the **Name** field, enter the first Cluster Member name. Example: `member1` c. In the **IPv4 address** field: If you manage the cluster from the same Virtual Network, enter the Cluster Member's private IP address. Otherwise, enter the Cluster Member's public IP address. d. Click **Communication**. e. In the **One-Time Password** field, enter the SIC key you set up in Azure. f. In the **One-Time Password** field, enter the SIC key again. g. Click **Initialize**. If the **One-time-password** is confirmed, the **Trust State** field shows Trust Established. h. To close the **Communication** properties window, click **Close**. If the Activation Key is confirmed, the **Trust State** field shows **Trust Established**. i. Click **OK**.
6	Repeat the Step 6 to add the second Cluster Member.
7	Click **Network Management** > **Get Interfaces** > **Get Interfaces With Topology**. If this warning appears: “Topology and Anti-Spoofing settings that are already defined will be overwritten. by results of this operation that contradict them, if any.Do you want to continue?” Click **Yes**.
8	Configure the interfaces **eth0** and **eth1**. a. Double-click the interface **eth0**. The **Network eth0** window shows. b. From the **General** tab, in the **Network type** field, select **Cluster + Sync**. c. In the **Virtual IPv4** field, enter the private VIP address and subnet mask of the cluster. In the diagram, the private VIP address is: 10.0.1.6 **Note -** You can find the cluster private VIP address in the Azure portal when you select the Active Cluster Member primary **NIC** > **IP configuration** > **"cluster-vip"**. d. From the **Network eth0** window, click **Topology** and disable the Anti-Spoofing. e. Click **OK**. f. Double-click the interface **eth1**. The **Network eth1** window shows. g. From the **General** tab, in the **Network type** field, select **Sync**. h. From the **Network eth1** window, click **Topology** and disable Anti-Spoofing. i. Click **OK**.
9	Install the applicable Access Control Policy on the cluster object.

Step 6: Configure NAT Rules

Note - See *Creating Objects in SmartConsole*.

In SmartConsole, create the NAT rules below to provide Internet connectivity from the internal subnets:

No	Original Source	Original Destination	Original Services	Translated Source	Translated Destination	Translated Services	Install On	Comments
1	Virtual Network	Virtual Network	*Any	= Original	= Original	= Original	Cluster object	Avoid NAT in the Virtual Network
2	App-subnet	App-subnet	*Any	= Original	= Original	= Original	Cluster object	
3	App-subnet	*Any	*Any	App-subnet (hidden address)	= Original	= Original	Cluster object	
4	Web-subnet	Web-subnet	*Any	= Original	= Original	= Original	Cluster object	
5	Web-subnet	*Any	*Any	Web-subnet (hidden address)	= Original	= Original	Cluster object	

Notes about the NAT rules:

- Rule 1 - You have to define this NAT rule *manually*.
- Rules 2 - 5 - SmartConsole creates these NAT rules *automatically*.
- Traffic between the *Web-subnet* and the *App-subnet* is based on the UDR rules. Each subnet has its own routing table.

For each internal subnet, create a network object:

Step	Description
1	Double-click the Web-subnet object. The **Web-subnet object** window shows.
2	Select the **NAT** tab > **Add automatic address translation rules**.

Step	Description
3	In the **Translation method** field, select **Hide** > **Hide Behind Gateway**.
4	In the **Install on Gateway** field, select the cluster object.
5	Click **OK**. This creates the *automatic* NAT rules.
6	Install the applicable Access Control Policy on the cluster object.

Step 7: Set Up the External Load Balancer in Azure

By default, the template you deploy creates an External Load Balancer, with the name `frontend-lb`, which faces the Internet.

The External Load Balancer sends health probes to TCP port 8117 to determine the health of the CloudGuard IaaS Security Gateways.

Create the load balancing rules in the Azure portal to allow incoming connections:

1. Go to **External Load Balancer > Frontend IP configuration**.
2. Click **Add**.

Notes:

- You cannot use these ports for forwarded traffic:
  - 80
  - 443
  - 444
  - 8082
  - 8080
  - 8117
- Do not change the health probe port.
- The Check Point cluster resource group includes an NSG associated with the frontend subnet. By default, the NSG allows all outbound and inbound traffic.
- The Load Balancer can be set up to listen on additional ports or on additional public IP addresses.

For more information, see *Multiple Frontends for Azure Load Balancer*. For an example, go to *"Step 9: Configure the Load Balancer to Listen on Multiple IP Addresses in Azure" on the next page*.

Step 8: Create Dynamic Object LocalGatewayExternal in SmartConsole

In SmartConsole, create the Dynamic object called *LocalGatewayExtenal*.

This object represents the private Cluster Member's IP addresses.

You use this Dynamic object in the next step.

1. In SmartConsole, click the **Objects** menu > **Object Explorer**.
2. From the top toolbar, click **New** > **Network Object** > **Dynamic Objects** > **Dynamic Object**.
3. In the **Enter Object Name** field, enter (case-sensitive):

   `LocalGatewayExternal`
4. Click **OK**.
5. Close the **Object Explorer**.

Step 9: Configure the Load Balancer to Listen on Multiple IP Addresses in Azure

Configure the Load Balancer to listen on additional public IP addresses. This setup is useful if you want the Security Gateway to secure multiple web applications, each with its own public IP address.

Configure the Load Balancer to listen on a second public IP address on TCP port 80, and then forward the traffic to the Check Point CloudGuard Security Gateway to TCP port 8083.

To configure the frontend pool:

Step	Description
1	In the Azure portal, select the Load Balancer called `frontend-lb`. **Note** - The Load Balancer is in the resource group you created.
2	Allocate a new public IP address: a. Click **Frontend IP configuration**> **Add**. b. Select a **Name**. **Example:** `<cluster>-app-2` c. Select the public **IP address** you created. d. Click **OK**.

Step	Description
3	Add a load balancing rule. a. Click **Load balancing rules** > **Add**. b. Enter the rule name. **Example:** `<cluster>-app-2-tcp-80` c. In the **Frontend IP address** field, select the newly created Frontend IP address. d. In the **Protocol** field, select **TCP**. e. In the **Port** field, enter 80. f. In the **Backend port** field, enter 8083. g. In the next **Backend pool** field, select the pre-existing cluster pool. h. In the **Health probe** field, select the health probe created by default by the template (TCP, port 8117). i. In the **Session persistence** field, select **None**. j. Set the desired **Idle timeout**, in minutes. k. In the **Floating IP** field, select **Disabled**. l. Click **OK**.

Configure Access Control Rule in SmartConsole

In SmartConsole, create a corresponding Access Control rule for External Load Balancer with these values:

Rule Name	Meaning / Value
Rule No	1
Name	Desired rule name
Source	`*Any`
Destination	`LocalGatewayExternal`
VPN	`*Any`
Service and Applications	The service object that represents the internal port
Data	`*Any`
Action	`Accept`
Track	`Log`
Install On	`*Policy Targets`

Load Balancer Conditions

The Active Cluster Member uses NAT to forward traffic that belongs to the two web applications, to the right web server.

NAT rules are defined with the special Dynamic Object.

The Dynamic object `LocalGatewayExternal` represents the private IP addresses of the external interface of Member 1 and Member 2.

For more information, see *"Step 8: Create Dynamic Object LocalGatewayExternal in SmartConsole" on page 26*.

No	Original Source	Original Destination	Original Services	Translated Source	Translated Destination	Translated Services	Install On
1	*Any	LocalGatewayExternal	TCP 8081	= Original	s App	https	Policy Targets
2	*Any	LocalGatewayExternal	TCP 8083	= Original	s Web	http	Policy Targets

Step 10: Configure VPN

In SmartConsole, create a Network Group object to represent the encryption domain for the cluster.

To create an object for the VPN configuration, see *Creating Objects in SmartConsole*.

For more information, see the *Check Point Security Management Administration Guide* for your Management Server version.

Step	Description
1	Create a Network Group object to represent the encryption domain of the cluster: a. In SmartConsole, click the **Objects** menu > **Object Explorer**. b. From the top toolbar, click **New** > **Network Group**. c. In the **Enter Object Name** field, enter the desired name. d. Click the **+** icon and select the applicable network objects. e. Click **OK**. f. Close the **Object Explorer**.
2	Edit the cluster object: a. In SmartConsole, from the left navigation panel, click **Gateways & Servers**. b. Double-click the cluster object. The **Gateway Cluster Properties** window shows.

Step	Description
3	Define your Network Group as the encryption domain of the cluster object: a. In SmartConsole, from the left navigation panel, click **Gateways & Servers**. b. Double-click the cluster object. The **Gateway Cluster Properties** window shows. c. In the cluster object left tree, click **Network Management** > **VPN Domain.** d. Select **Manually defined**. e. In the right corner of this field, click the **[...]** button and select the Network Group object you created in Step 1.
4	Define the VPN community: a. In the cluster object left tree, click **IPsec VPN**. b. In the section **This Security Gateway participates in the following VPN Communities**, select the applicable VPN community.
5	Define the outgoing VPN interface: a. In the cluster object left tree, click **IPsec VPN** > **Link Selection**. b. In the **IP Selection by Remote Peer** section, select **Always use this IP address** > **Main address**. c. In the **Outgoing Route Selection** section: i. Click **Source IP address settings**. ii. Select **Manual**. iii. Choose **Selected address from topology table.** iv. Select the private cluster object VIP address. v. Click **OK**. d. In the **Tracking** section, select the desired option. e. Click **OK** to close the **Gateway Cluster Properties** window.
6	Configure the VPN Community to use Permanent Tunnels: a. In SmartConsole, click the **Objects** menu > **Object Explorer**. b. In the left tree, clear all boxes except for **VPN Communities**. c. Double-click the VPN community, in which this cluster object participates. The **VPN Community** window shows. d. In the left tree, click **Tunnel Management**. e. Select **Set Permanent Tunnels**. f. Select the applicable option. g. Click **OK** to close the VPN Community properties window. h. Close the **Object Explorer**.
7	Install the applicable Access Control Policy on the cluster object.

Additional Information

Testing and Troubleshooting

You can use the APIs to retrieve information about the cluster resource group.

Use these commands on each Cluster Member to confirm that the cluster operates correctly:

Run these commands in the Expert mode:

```
cphaprob state
cphaprob -a if
```

Example:

```
[Expert@HostName:0]# cphaprob state
Cluster Mode: High Availability (Active Up) with IGMP Membership
Number Unique Address Assigned Load State
1 (local) 10.0.1.10 0% Active
2 10.0.1.20 100% Standby
```

Use the cluster configuration test script on each Cluster Member to confirm it is configured correctly:

The script verifies:

- The configuration is defined in the `$FWDIR/conf/azure-ha.json` file, which is created by the ARM template.
- A Primary DNS server is configured and works.
- The machine is set up as a Cluster Member.
- IP forwarding is enabled on all network interfaces of the Cluster Member.
- It is possible to use the APIs to retrieve information about the cluster's resource group.
- It is possible to log in to Azure with the Azure credentials in the `$FWDIR/conf/azure-ha.json` file.
- Calibration of ClusterXL configuration for Azure. cluster

To get the latest version of the test script:

Important - In a Cluster, you must configure all the Cluster Members in the same way.

1. Download the latest version of the test script.

- For R80.40 and above, use this link.
- For R80.30 image version R80.30.273.590 and above, use this link.
- For other images, use this link.

Note - To get the image version, see sk116585.

2. Copy the downloaded script to a directory.
3. Connect to the command line and log in to the Expert mode.
4. Back up the current `$FWDIR/scripts/azure_ha_test.py` script:

```
cp -v $FWDIR/scripts/azure_ha_test.py{,_backup}
```

5. Copy the latest script to the `$FWDIR/scripts/` directory:

```
cp -v /<path to the downloaded script package>/azure_ha_test.py
$FWDIR/scripts/
```

6. Assign the required permissions:

```
chmod -v 755 $FWDIR/scripts/azure_ha_test.py
```

To run the script on each Cluster Member:

1. Connect to the command line.
2. Log in to the Expert mode.
3. Run the script with this command (do not change the syntax):

```
$FWDIR/scripts/azure_ha_test.py
```

If all tests were successful, this message appears:

```
All tests were successful!
```

Otherwise, an error message appears with information to troubleshoot the problem.

A list of common configuration errors:

Message	Recommendation
`The attribute (ATTRIBUTE) is missing in the configuration`	
`Primary DNS server is not configured` `Failed to resolve (host)`	The Cluster Member is not configured with a DNS server.
`Failed in DNS resolving test`	Confirm that DNS resolution on the Cluster Member works.

Message	Recommendation
`You do not seem to have a valid cluster configuration`	Make sure that the Cluster Member configuration on the Check Point Security Management Server is complete and that the Security Policy is installed.
`IP forwarding is not enabled on Interface (Interface-name)`	Use PowerShell to enable IP forwarding on all the network interfaces of the Cluster Member.
`failed to read configuration file: /opt/CPsuite-R80/fw1/conf/azure-ha.json`	The Azure Cluster Member configuration is not up to date, or written correctly.
`Testing credentials`	Failed to log in with the credentials provided. See the exception text to understand why.
`Testing authorization (Exception)`	Make sure the Azure Active Directory service account you created is designated as a Contributor to the cluster resource group.

Simulate a cluster failover:

For example, shut down the internal interface of the Active Cluster Member.

1. On the current Active Cluster Member, run in the Expert mode:

   ```
   ip link set dev eth1 down
   ```

2. In a few seconds, the second Cluster Member has to report itself as the Active Cluster Member.

   Examine the cluster state on each Cluster Member in the Expert mode:

   ```
   cphaprob state
   ```

3. On the former Active Cluster Member, run in the Expert mode:

   ```
   ip link set dev eth1 up
   ```

If you experience issues:

- Make sure you have a configured Azure Active Directory Service Account.

  The service has to have:

  - Contributor privileges to the resource group
  - At least minimum privileges on the Cluster Member deployment resources. See *"Changing Template Components" on page 36*.

- To make the networking changes automatically, the Cluster Members have to communicate with Azure. This requires HTTPS connections over TCP port 443 to the Azure end points. Make sure the Security Policy that is installed on the Cluster Members allows this type of communication.

Using the Azure High Availability Daemon

The cluster solution in Azure uses the daemon to make API calls to Azure when a cluster failover takes place.

This daemon uses a configuration file `$FWDIR/conf/azure-ha.json` on each Cluster Member.

When you deploy the solution above from the template supplied, this file is created automatically.

The configuration file is in JSON format and contains these attributes:

Attribute name	Type	Value
`debug`	Boolean	`true` or `false`
`subscriptionId`	String	Subscription ID.
`location`	String	Resource group location.
`environment`	String	Name of the environment.
`resourceGroup`	String	Resource group name.
`credentials`	String	IAM. Indicates using automatic credentials on the Cluster Member Virtual Machine.
`proxy`	String	Name of the proxy.
`virtualNetwork`	String	Name of the Virtual Network.
`clusterName`	String	Name of the cluster.
`templateName`	String	Name of the template.
`tenantId`	String	ID of the tenant.

Note - If you use your own service principal, the `credentials` attribute contains:

- Your Client-ID
- Your Client-secret
- Grant type `client-credentials`
- Your Tenant ID

You can confirm that the daemon in charge of communicating with Azure runs on each Cluster Member.

From the Expert mode, run:

```
cpwd_admin list | grep -E "PID|AZURE_HAD"
```

The output should look like in this example:

```
APP         PID    STAT   #START   START_TIME              MON   COMMAND
AZURE_HAD   3663   E      1        [12:58:48] 15/1/2016    N     python
/opt/CPsuite-R80.20/fw1/scripts/azure_had.py
```

Notes:

- The script appears in the output:
  - The `STAT` column should show **E** (executing)
  - The `#START` column should show **1** (the number of times this script was started by the Check Point WatchDog)

To troubleshoot issues related to this daemon, generate debug. From the Expert mode, run:

- To enable debug printouts:

  ```
  azure-ha-conf --debug --force
  ```

- To disable debug printouts:

  ```
  azure-ha-conf --no-debug --force
  ```

The debug output is written to `$FWDIR/log/azure_had.elg*` files.

Using a Different Azure Cloud Environment

If you want to deploy your cluster in an environment other than the standard Azure environment, make sure to edit this file:

`$FWDIR/conf/azure-ha.json`

Example:

```
{
...
    "environment": "[Azure-cloud-environment]",
...
}
```

The Azure-Cloud-Environment has to be one of these:

- Azure Cloud (the default global cloud environment)
- Azure China Cloud
- Azure US Government
- Azure German Cloud

Procedure:

1. From the Expert mode, run:

   ```
   azure-ha-conf --environment '<Azure-cloud-environment>' --force
   ```

2. Make sure the file syntax is correct. From the Expert mode, run:

   ```
   python -m json.tool $FWDIR/conf/azure-ha.json
   ```

3. Apply the changes. From the Expert mode, run:

   ```
   $FWDIR/scripts/azure_ha_cli.py reconf
   ```

   Note - If you deploy in the default global cloud environment, you can omit this attribute.

Important note about the service principal:

If you use any of these different environments, you have to create your own service principal. No default service principal is created.

Working with a Proxy

In some deployments, you can only access the Internet through a web proxy.

To allow the Cluster Member to make API calls to Azure through the proxy, edit the `$FWDIR/conf/azure-ha.json` file and add this attribute:

```
{
...
    "proxy": "http://[Proxy-Server]:[Proxy-Port]",
...
}
```

- *Proxy-Server* is the host name or IP address of the web proxy server
- *Proxy-Port* is the port on the proxy server

Note - The URL scheme has to be HTTP and not HTTPS.

Example:

```
{
...
    "proxy": "http://proxy.example.com:8080",
...
}
```

Procedure:

1. Change the proxy settings. From the Expert mode, run:

   ```
   azure-ha-conf --proxy 'http://[Proxy-Server]:[Proxy-Port]' --force
   ```

2. Make sure the file syntax is correct. From the Expert mode, run:

   ```
   python -m json.tool $FWDIR/conf/azure-ha.json
   ```

3. Apply the changes. From the Expert mode, run:

   ```
   $FWDIR/scripts/azure_ha_cli.py reconf
   ```

Changing Template Components

The Check Point cluster's public IP address has to be in the same resource group as the Cluster Members.

These resources can be in any resource group:

- Virtual Network
- Network interfaces
- Route tables
- Storage account

Note - Make sure the resources Virtual Network and External Network Interfaces use the same automatic service principal with the same permissions.

Naming Constraints

Do not change the name of any resources.

Cluster Members VM names must match the Cluster name with a suffix of '1' and '2'.

Example:

```
<member_name1>
<member_name2>
```

Network Interface names must match the Cluster Member VM names with a suffix of '-eth0' and '-eth1'.

Example:

```
<member_name1-eth0>
<member_name1-eth1>
<member_name2-eth0>
<member_name2-eth1>
```

The IP address of the cluster has to match the configuration file.

By default it should match the cluster name.

Permissions

It is possible to assign service principal permissions to specific Azure resources. See sk116585 for information on how to find the image version.

To allow the cluster to update the necessary Azure resources on failover, the service principal has to be assigned at least these roles on these resources or on their respective resource group:

Resource Type	Role
Any public IP address attached to the External Load Balancer	Virtual Machine contributor
Public Load Balancer	Network contributor
CloudGuard Virtual Machines	Reader
Cluster public IP address	Network contributor
Public IP address of each Cluster Member	Virtual Machine contributor
Virtual Network	Virtual Machine contributor
The external network interfaces (`eth0`) used by the Cluster Member	Virtual Machine contributor

Creating Objects in SmartConsole

For more information, see the *Check Point Security Management Administration Guide* for your Management Server version.

Important - After you create an object, you must publish the session to save the changes in the management database.

To create a Host object:

1. From the top right **Objects Pane**, click **New** > **Host**.

   The **New Host** window shows.
2. In the **Machine** field, enter the private IP address of the machine.

To create a Network object:

1. From the top right **Objects Pane**, click **New** > **Network**.

   The **New Network** window opens.
2. Enter the **Object Name** (specifically the subnet name).
3. Enter the **Network address** and **Net mask**.

To create a Service (port) object:

1. From the top right **Objects Pane**, click **New** > **More** > **Service**.
2. Select your TCP/UDP service.
3. Enter the **Object name**.

4. In the **Enter Object Comment** field, enter the port name.
5. In the **General** field, select your **Protocol**.
6. In the **Match By** field, select the **Port** number.
7. Click **OK**.

To create a Network Group object:

1. From the top right **Objects Pane**, click **New** > **Network Group**.

   The **New Network Group** window opens.
2. Click **+** to select your internal subnets.
3. Click **OK**.

Related Solutions

- sk109360 - Check Point Reference Architecture for Azure
- sk113583 - How to add a network interface to a Check Point Security Gateway in Azure
- sk113476 - Azure Virtual Network peering

Upgrading a Check Point CloudGuard IaaS High Availability Solution to a Newer Version

Use these instructions to upgrade a deployed Check Point CloudGuard IaaS High Availability.

Notes:

- The upgrade maintains the network configurations used in the existing Check Point CloudGuard IaaS High Availability solution.
- The Cluster VIP resource name is *cluster object name* defined in deploying the solution.

Step	Description
1	Log in to the Azure portal.
2	Deploy a new Check Point CloudGuard IaaS High Availability in the needed version. You need to deploy it to the same subnets as in the existing CloudGuard IaaS High Availability solution.
3	Provide an access role with Contributor permission for the new CloudGuard IaaS High Availability members to these resource groups: 1. The resource group that contains the original VNet. 2. The resource group that contains the old cluster IP addresses and the Load Balancer. For more information, see *"Step 2: Set Credentials in Azure" on page 17*. **Note** - Access is granted after a few minutes to an hour.
4	Setup routes on Cluster Members to the Internal Subnets. See *"Step 4: Set Up Routes on Cluster Members to the Internal Subnets " on page 21*.

Step	Description

5	Use the old cluster VIP address:
	* On each Cluster Member, edit the configuration file located at `$FWDIR/conf/azure-ha.json`
	Under `"clusterNetworkInterfaces" -> "etho"`, replace the public IP name with the original cluster VIP resource id.
	Example:
	■ Old:
	`"clusterNetworkInterfaces": { "eth0": [ "10.72.0.6", "newHAClusterIp" ] },`
	■ New:
	`"clusterNetworkInterfaces": { "eth0": [ "10.72.0.6", "/subscriptions/123/resourceGroups/ha-rg/providers/Microsoft.Network/publicIPAddresses/originalHAClusterIp" ] },`

Step	Description
6	Update the existing cluster object in SmartConsole: **Important** - Do not install policy. a. In SmartConsole, double-click on the existing cluster object. b. In **Cluster Members**, update each member to match the compatible member of the new Check Point CloudGuard IaaS High Availability. Enter the **IPv4 address**, and then create a SIC connection (after resetting the current communication). If you are managing the cluster from the same Virtual Network, then enter the Cluster Member's private IP address. Otherwise, enter the Cluster Member's public IP address. c. In **General Properties** > **Platform**, update the version of your new CloudGuard IaaS High Availability, and then click **Get**. d. In **Network Management**: i. Double-click the interface **eth0**: In the **Virtual IPv4** field, and then enter the private VIP address and subnet mask of the new CloudGuard IaaS High Availability. ii. For both **eth0** and **eth1**, modify **the Members IPs** to match the new CloudGuard IaaS High Availability members IP addresses. Also, enter the external private IPs in eth0 and the internal private IPs in eth1. e. For a VPN configuration in **IPsec VPN**, select **Link selection**. In the **Outgoing Route Selection**: i. Click **Source IP address settings**. ii. Select **Manual**. iii. Choose **Selected addresses from topology table**. iv. Select the private cluster object VIP address of your new CloudGuard IaaS High Availability for Azure.
7	Delete the External Load Balancer, the Internal Load Balancer, and the public cluster address created in the deployment of the new Check Point CloudGuard IaaS High Availability (step 1). They are located in the new Check Point CloudGuard IaaS High Availability resource group.
	 **Important** - Connectivity will be lost during the next steps.
8	Add the new Check Point CloudGuard IaaS High Availability's members to the backend pools: For each Load Balancer used in the original solution, add the new members to the existing backend pools. Make sure to select the right IP address (private internal for the backend Load Balancer and private external to the frontend Load Balancer).
9	Delete the old Check Point CloudGuard IaaS High Availability's members from the backend pools of each Load Balancer used in the original solution.
10	Install the applicable Policy on the cluster object.

Step	Description
11	Detach the cluster VIP from the original Check Point CloudGuard IaaS High Availability's members: a. Stop both original Check Point CloudGuard IaaS High Availability's members. b. Select the **Active Cluster Member's primary NIC** > **IP configuration** > **"cluster-vip"** in the original Check Point CloudGuard IaaS Cluster and delete it. Initiate a failover in the new CloudGuard IaaSHigh Availability to attach the original cluster VIP to the new members.
	Note - After an access role to the new members has been granted, the new CloudGuard IaaS High Availability now handles all traffic in the environment (inbound, outbound, E-W, VPN Tunneling). Verify that all the traffic flows work as expected before proceeding, and that the cluster failover works as expected.
12	Delete the original Check Point CloudGuard IaaS Cluster and other redundant resources. **Note -** If you are using resources from the old resource group, such as VNETs or cluster IP addresses, do not delete them.

Upgrade a Check Point Cluster to the CloudGuard IaaS High Availability Solution

Use these instructions to upgrade a deployed Check Point CloudGuard IaaS Cluster for the CloudGuard IaaS High Availability solution.

Notes:

- All public IP addresses associated with the Cluster solution will change. This is due to the fact that a "Standard" SKU Load Balancer is used in the High Availability solution. "Basic" SKU public IP address resources cannot be associated with "Standard" SKU Load Balancers.
- All VPN endpoints must be updated to use a new public VIP.
- Before starting the upgrade, read the steps below and prepare an upgrade plan.

Step	Description
	Important - Before you begin, verify that your deployed Management Server or Multi-Domain Server is able to manage the version of the CloudGuard IaaS High Availability solution that you intend to upgrade to. If the Management Server is an earlier version, you may need to install a Hotfix. Contact Check Point Support for more information.
1	Log in to the Azure portal.
2	Deploy a new Check Point CloudGuard IaaS High Availability solution. **Use the same network configurations as in the existing CloudGuard IaaS Cluster solution**. **Important** - When deploying an existing VNet, you must either create or modify the Network Security Groups, since they are not automatically created.
3	Provide an access role for the virtual network in which the new CloudGuard IaaS High Availability solution is deployed to the CloudGuard IaaS High Availability members. For more information see *"Step 2: Set Credentials in Azure" on page 17*. **Note** - Granting access will take a few minutes to an hour.
4	Setup routes on Cluster Members to the Internal Subnets. See *"Step 4: Set Up Routes on Cluster Members to the Internal Subnets " on page 21*
5	Create load balancing rules: ■ For each NAT rule in the Cluster solution's **External Load Balancer**, create a compatible load balancing rule for the new CloudGuard IaaS High Availability's External Load Balancer.

Step	Description
6	Update your existing cluster object in SmartConsole: **Important** - Do not install policy. 1. In SmartConsole, double-click on the existing cluster object. 2. In **General Properties**: a. In the Virtual IPv4 address, enter the public address allocated for the new CloudGuard IaaS High Availability. **Note -** The cluster IP address is found in the Azure portal by selecting the Active Cluster Member's **NIC** > **IP configuration** > "**cluster-vip**". b. If needed, update the version. 3. Under "**Cluster Members**" update each member to match the compatible member of the new Check Point CloudGuard IaaS High Availability. Enter the **IPv4 address** and then create a SIC connection (after resetting the current communication). If you manage the cluster from the same Virtual Network, enter the Cluster Member's private IP address. Otherwise, enter the Cluster Member's public IP address. 4. In **Network Management**: a. Double-click the interface **eth0**: In the **Virtual IPv4** field, and then enter the private VIP address and subnet mask of the new CloudGuard IaaS High Availability. b. For both **eth0** and **eth1**, **Modify** the "**Members IPs**" to match the new IP addresses of CloudGuard IaaS High Availability members. Enter the external private IPs in eth0 and internal private IPs in eth1. 5. For a VPN configuration, click **IPsec VPN** > **Link selection**. In the **Outgoing Route Selection**: a. Click **Source IP address settings**. b. Select **Manual**. c. Click **Selected address from topology table**. d. Select the private cluster VIP address of the new CloudGuard IaaS High Availability.
7	In SmartConsole, configure the policy and NAT rules as needed  **Important** - Connectivity will be lost during the next steps.
8	Update route tables in the azure portal: ■ For each route in the App or Web Internal subnet where the **Nexthop address** matches the active member of your old Check Point cluster, change it to match the iLB-internal-address. **Note** - If the subnet houses the Security Management Server, managing the Cluster Members, then do not change the route which allows the Security Management Server to communicate directly with each Cluster Member. For more information see "*"Step 3: Set Up Internal Subnets and Route Tables" on page 19*. ■ In the frontend and backend subnets (in places that **Nexthop address** match the active member of your old Check Point cluster), change it to match the iLB-internal-address.
9	Install the applicable Policy on the cluster object.

Step	Description
	Note - After an access role to the new members has been granted, the new CloudGuard IaaS High Availability now handles all traffic in the environment (inbound, outbound, E-W, VPN Tunneling). Verify that all the traffic flows work as expected before proceeding
10	Delete the original Check Point CloudGuard IaaS Cluster and other redundant resources. **Note** - If you are using resources from the old resource group, such as VNets, do not delete them.

Known Limitations

- Support for Jumbo Hotfix Accumulators:
  - See sk109141 for information about supported Jumbo Hotfix Accumulators.
- Only two Cluster Members in a cluster are supported.
- Only High Availability Mode (Active/Standby) is supported. Load Sharing Mode is not supported.
- VRRP cluster is not supported.
- Only the Active Cluster Member can reach services from the cluster through VPN.

  The Standby Cluster Member can reach those services only when it becomes the Active Cluster Member.
- For outbound and VPN traffic, you cannot delete or disable the public IP addresses of Cluster Members.
- The feature is only available in Azure Resource Manager deployments.

  It is not supported with Azure Service Manager (also known as classic) deployments.
- When you use the standard Internal Load Balancer it does not support Stateful failover.
- Managed service identity for Virtual Machines is only available in the Azure Cloud environment. Other environments require a manual service identity management.

Terms

Active

State of a Cluster Member that handles network connections that pass through the cluster. In a cluster deployment, only one Cluster Member is Active and can handle connections.

Active Directory (AD)

Microsoft® directory information service. Stores data about user, computer, and service identities for authentication and access.

Availability Set

A collection of Virtual Machines that are managed together to provide application redundancy and reliability. The use of an availability set ensures that during either a planned or unplanned maintenance event at least one Virtual Machine is available. (Description from the Microsoft Azure glossary).

Azure Environment

An Azure environment an independent deployment of Microsoft Azure, such as Azure Cloud for global Azure and Azure China Cloud for Azure operated by 21Vianet in China.

Azure PowerShell

A command-line interface to manage Azure services via a command line from Windows. (Description from the Microsoft Azure glossary)

Check Point WatchDog

A process that launches and monitors critical processes such as **Check Point** daemons on the local machine, and attempts to restart them if they fail.

Cluster

Two or more Security Gateways that work together in a redundant configuration - High Availability.

Failover

Also, Fail-over. Transferring of a control over traffic (packet filtering) from a Cluster Member that suffered a failure to another Cluster Member (based on internal cluster algorithms).

Load Balancer

A resource that distributes incoming traffic among computers in a network. In Azure, a Load Balancer distributes traffic to Virtual Machines defined in a Load Balancer set. A Load Balancer can be Internet-facing, or it can be internal. (Description from the Microsoft Azure glossary)

Resource

An item that is part of your Azure solution. Each Azure service enables you to deploy different types of resources, such as databases or Virtual Machines. (Description from the Microsoft Azure glossary)

Resource Group

A container in Resource Manager that holds related resources for an application. The resource group can include all of the resources for an application, or only those resources that are logically grouped together. You can decide how you want to allocate resources to resource groups based on what makes the most sense for your organization.

SmartConsole

Check Point main GUI client used to create and manage the Security Policy.

Standby

State of a Cluster Member that is ready to be promoted to Active state (if the current Active Cluster Member fails). Applies only to ClusterXL High Availability Mode.

Subnet

A logical subdivision of an IP network.

User Defined Routing

A route table or a set of rules to create network routes, so that yourVirtual Machine can handle the traffic between subnets and to the Internet.

Virtual Machine (VM)

The software implementation of a physical computer that runs an operating system. Multiple Virtual Machines can run simultaneously on the same hardware. In Azure, Virtual Machines are available in a variety of sizes. (Definition from the Microsoft Azure glossary).

Virtual Network

A network that provides connectivity between your Azure resources that is isolated from all other Azure tenants. An Azure VPN Gateway lets you establish connections between Virtual Networks and between a Virtual Network and an on-premises network. You can fully control the IP address blocks, DNS settings, Security Policies, and route tables within a Virtual Network. (Description from the Microsoft Azure glossary)

Check Point Copyright Notice